Virtue in Context

by

Andrew Charles Ball

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Philosophy

University of Alberta

Abstract

Virtue Reliabilism and Virtue Responsibilism are two theories within the enterprise of Virtue Epistemology. The former considers virtues to be those competences whose reliability is what confers justification on its product beliefs. The latter considers virtues as being those deep-seated intellectual traits that are part of a person's very character, and so when such virtues are possessed and exercised by an agent, they achieve beliefs that are justified via being the products of virtue.

Both theories face difficult objections, however. Virtue reliabilism is challenged by the generality problem which claims that since justification is determined by how reliable the belief-forming process is, what we have to do to figure out just how reliable a process is will be by coming up with a proper description of the relevant process. However, there is no principled way to come up with such descriptions in order to determine the correct level of generality regarding the description of a belief-forming process. Virtue responsibilism is challenged by the situationist critique which claims that virtues do not do the kind of work we think they do. Rather, when agents morally and intellectually act in praiseworthy ways, their success is due to situational features of their experience and not to anything like what we think of as virtue. What both problems have in common is that they require their respective target theories to come up with explanations that accurately describe

what's going on when people get things right epistemically. As such, virtue reliabilism and virtue responsibilism both face the same sort of problem.

In this dissertation, I attempt to provide a contextual response to the generality problem that doesn't solve that problem per se, but suggests that the problem itself is problematic. In requiring the reliabilist to come up with a universal principle(s) that allows the epistemic practitioner to determine the relevant description for reliable belief-forming processes, the proponent of the generality problem assumes that processes are relevantly alike in all instances of belief formation. But our epistemic practices show otherwise. Indeed, standards for belief justification vary between different situations and contexts. Some contexts have quite stringent standards regarding how reliable a process must be in order to be considered reliable enough while other contexts are not so stringent. As such, that there could be some universal standard or principle for process description is impossible.

This dissertation argues for a contextual resolution to the generality problem. Furthermore, it argues that the Virtue Reliabilism of Ernest Sosa seems to have embedded in it a very congenial affinity to attributor contextualism. In light of (i) those arguments, (ii) the similar contextual nature of neo-Aristotelian aretaic thought, (iii) as well as the similarity of the situationist critique to the generality problem, a similar contextual response to that critique is then offered.

Preface

This thesis is an original work by Andrew Charles Ball. No part of this thesis has been previously published in any form.

For ASHLEY,
my Love,
whose affection, encouragement, and prayers have touched every page herein.

Acknowledgements

Thanks to my committee: my supervisor Bruce Hunter for his invaluable advice, support, and direction throughout the writing of this thesis; Jennifer Welchman and Ingo Brigandt for their insight and suggestions. I also want to thank Thomas Spalding (Dept. of Psychology) and Jeremy Fantl (University of Calgary, Dept. of Philosophy) for the time they took to read this thesis and offer very valuable suggestions.

Thanks also to those FINANCIAL CONTRIBUTORS who have made it possible that I could get this far in my education without incurring any debt; the University of Detroit Mercy; the United Auto Workers; the DaimlerChrysler Corporation Foundation; the Faculty of Graduate Studies at the University of Windsor; the Centre for Research in Reasoning, Argumentation, and Rhetoric at the University of Windsor; the Faculty of Gradate Studies and Research and the Department of Philosophy at the University of Alberta.

And finally, a special thanks to those whose love and friendship have been an invaluable strength and encouragement to me through my graduate school career: Ashley; my parents, Richard and Kimberly Ball; my in-laws, David and Shawna Schultz; Matt Pinson; Tom McCullough; Harry Veryser; Tyler Hanson; Tom Chacko; John Wood; Paul Tan; Kervin Raugust; the folks of Knox Evangelical Free Church in Edmonton and Emmanuel Evangelical Missionary Church in Holden.

TABLE OF CONTENTS

Four: Virtue and Context in Responsibilist Epistemology

Five: The Generality Problem and the Situationist Critique

Six: A Contextual Resolution

Introduction

> *Do some epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them, just as I told you to, you are no longer ignoring them, properly or otherwise. So you have landed in a context with an enormously rich domain of potential counter-examples to ascriptions of knowledge. In such an extraordinary context, with such a rich domain, it never can happen (well, hardly ever) that an ascription of knowledge is true. Not an ascription of knowledge to yourself (either to your present self or to your earlier self, untainted by epistemology); and not an ascription of knowledge to others. That is how epistemology destroys knowledge… that is how knowledge is elusive. Examine it, and straightway it vanishes."*
>
> – David Lewis (1996, p. 559)

Virtue has an interesting history. Since the time of Plato, it has often been appealed to as the explanatory factor for good moral behavior. More recently (but going back to the Greeks as well),[1] there has been a growing movement in epistemology that appeals to epistemic or intellectual virtue as the explanatory factor for good epistemic behavior understood as intellectual performances that result in knowledge. The development arising out of this latter phenomenon has come to be known as *virtue epistemology*. It, however, consists of two distinct strains. One is *virtue reliabilism*. This theory posits that

[1] Of course, Aristotle started this. But the current trend of philosophical development suggests a sort of revival of this work.

epistemic virtues are competences or abilities that agents have which form true beliefs. Because these competences are reliable at doing this (after all, they wouldn't be considered competences if they were unreliable), it is argued that their product beliefs are justified via their creator-virtue's reliability. The second strain is *virtue responsibilism*. It posits that intellectual virtues aren't mere competence or abilities, but deep-seated intellectual traits that are cultivated by a person out of their motivation for truth. As such, they have a strong and intimate connection to one's own character.

My own preference has been for virtue reliabilism. Why? Well first, I like talking about epistemic justification in terms of 'reliability' because it plausibly allows me to include all sorts of true beliefs as instances of knowledge than has traditionally been the case since the traditional conception of belief justification has required the agent to be cognitively aware of the reasons for her belief. Virtue reliabilism, however, doesn't require that of an agent. But secondly, I like how the qualification 'virtue' constrains and restrains 'reliability' as the justificatory factor by focusing on the epistemic competency and skill of the agent herself. That an epistemic process is virtue-reliable is tantamount to saying that the agent herself is epistemically reliable, and not by accident, but because she has aptly employed her epistemic competences with the aim of making true beliefs. Exercising her competences correctly and successfully not only make her beliefs true, but also makes them justified, and hence, knowledge.

But this isn't a perfect way to think about knowledge. Even if we like reliability, we have to admit that it is a less or more thing – it comes and goes in degrees. So, a belief produced by a virtuous process that is more reliable is, by definition, more justified (and vice versa). But here's the problem: how do we determine or discover the reliability of a true-belief-producing-process? Simply stated, this is accomplished through how the process gets described. But how are we to describe the process? If we go too broad in our description, we may unwittingly 'justify' beliefs that really aren't justified, and if we go too narrow, we might risk describing a process that could ever happen only once, and thus not repeatable, and thus not plausibly reliable (or unreliable). This is the infamous generality problem for reliabilism and says that our descriptions of belief forming processes need to convey the correct level of generality.

Is there an alternative virtue approach that doesn't fall prey to the generality problem? How about the more Aristotelian virtue responsibilism? On this approach, virtues are not processes *per se*, but deep entrenched character traits that cause one to be morally and intellectually successful. Applied to epistemology, cases of knowledge are true beliefs that have been attained as the result of an agent exercising those truth-conducive intellectual virtues that she has developed and thus exercises reliably (i.e., because the thought here is that she will habitually exercise them successfully). So whereas reliability is integrally tied to belief justification on the virtue-reliabilist account, reliability is the happy result of deep entrenched virtue on the virtue-

responsibilist account because, by definition, the virtuous person is consistently virtuous. Seemingly, then, the generality problem doesn't exist here.

However, there is a criticism that also gets its impetus by poking at the reliability claims of this approach. For Aristotle the morally virtuous person "always makes the best of circumstances" which implies that the virtuous person is perfectly reliable (1101a1-5). But as the criticism goes, this requires too much. John Doris and Mark Alfano, for instance, have argued that virtue doesn't really exist in the way this responsibilist approach says it does. If it did, we would be able to observe perfect reliability. We would see it expressed as "cross-situational consistency" (a sort of behavioral reliability), which means that no matter the situation virtue would *always* be exercised by the virtuous person and hence be the most relevant factor that explains why that person acted morally (Doris 1998, p. 507). But Doris and Alfano illustrate their claims by pointing to the results of social-psychological experiments where the virtue ought to be the explanatory factor of good moral and intellectual behavior, but it seemingly isn't. Their conclusion is that it is never the case that good actions are due solely to virtue, but that they include a large dose of what are seemingly irrelevant external or "situational" factors.

So, virtue responsibilism, while not vulnerable to the generality problem, still has a problem of its own – of explaining how intellectual virtue is supposed to be the explanation of why people behave epistemically well. But notice how similar the two problems are. They both seem to be taken by the

issue of how to describe good moral and intellectual action. With regard to how these problems apply epistemologically (which is what I am more interested in), they are instances of the problem of how to account for epistemic successes by specifying the sources of true beliefs in such a way that best serve the explanatory and normative interests we have for epistemic evaluation.

So how might we respond to these criticisms? In what follows in this dissertation, I want to seriously consider epistemic contextualism, looking through its lens in order to see what kind of solutions we might come up with. In chapter one, I delineate the attributor contextualism of Keith DeRose and Stuart Cohen, and I argue that out of their work a two-pronged contextualism is implied – one that understands there to be *micro*-contexts which are simply the situations in which we epistemically perform, and *macro*-contexts which are the social contexts that determine the rules/standards/norms for epistemic performance in the various micro-contexts.

In chapter two, I argue that Ernest Sosa's virtue reliabilism is a theory of epistemic justification that implicitly has an understanding of these contexts built into it. This is apparent in his argument that reliability is an epistemic desideratum that is of importance to the community, and so by communally determining what counts as being a reliable performance under (i.e., relative to) certain kinds of (different) conditions, we socially determine what it is that we want in people's epistemic performances.

But I also want to argue that virtue responsibilism has a similar contextual element as well. And, so, in chapter three I look at Aristotle and I argue that his *moral* virtue theory allows for a sort of neo-Aristotelian interpretation that finds macro and micro contexts implicitly relied upon given the way that Aristotle formulates it. Alasdair MacIntyre, then, is discussed as a neo-Aristotelian who also would subscribe to this sort of contextual overlay.

These are discussed because the epistemic virtue responsibilist theory I target as the model one in chapter four is the responsibilism of Linda Zagzebski who comes out of the Aristotelian tradition herself. As such, the kinds of claims and conclusions I make regarding Aristotle find a natural kinship to Zagzebski's theory as well. My conclusion there is that Zagzebski seems to imply that there is a macro-contextually determined epistemic standard (the virtuous exemplar) that determines the proper standards or norms for good epistemic behavior for the various micro-contexts one finds oneself in.

Chapter five is quite expository. There, I discuss the generality problem for virtue reliabilism and the situationist critique for virtue responsibilism. My conclusion is that they are similar enough so that the strategy that possibly resolves one might also plausibly resolve the other.

And, so, this is exactly what I try to do in chapter six. There, I argue that since reliability is a term that is so richly sensitive to context, the generality problem overreaches when it claims that reliabilism needs to be able

to come up with a principle that adequately/exhaustively describes the relevant belief-forming process for each situation in which a belief is formed. My argument is that every situation is different, not merely in terms of the facts about it, but also regarding what the standards for knowledge are in those situations. In some, reliability may mean something higher than 75%, for instance, while in others it may mean something much less. It all depends on *context* – what the *macro*-context deems important according to its needs/interests/purposes (i.e., stakes) in the various *micro*-contexts in which we epistemically perform. Hence, the generality problem is not a real problem as much as it's rooted in a misunderstanding of our actual epistemic practice.

Finally, then, this 'solution' is offered as a putative one to the situationist critique. As such, the critique may overreach in its claims that particular social-psychological experiments – which are designed to see whether virtue actually causes good behavior or if seemingly unrelated situational factors do this work – reveal virtue to be a farce. It may just be that responsibilist virtues don't always look the same to the naked eye in practice. And, of course, what any responsibilist would want is a theory that makes room for cooperation between virtues and situational factors. I argue that responsibilism does this and, as a result, the situationist critique misunderstands the broad role that virtue and practical wisdom play together.

Chapter One

Two Dimensions of Epistemic Contextualism

> *Since Kant, philosophers have hoped […] by finding the apriori structure of any possible inquiry, or language, or form of social life. If we give up this hope, […] we may gain a renewed sense of community. Our identification with our community – our society, our political tradition, our intellectual heritage – is heightened when we see this community as ours rather than nature's, shaped rather than found, one among many which men have made.*
>
> – Richard Rorty (1980, p. 727)

1.1. Getting our Bearings with Contextualism

I will always remember the first time I completely doubted whether I actually knew anything at all. This event happened in an undergraduate epistemology class. I was told about the (possible) evil genius that might (possibly) be the cause behind any proposition that I have ever claimed to know. I was told that by the simple probing of his doxastic-state-creating electro-wand on my vat-housed brain somewhere in a distant galaxy, he (possibly) has caused every belief that I have.[1] Furthermore, I was told that even though this (possible) explanatory story seemed either preposterous at best or a lame work of science fiction at worst, if I could not conclusively rule

[1] Here I take liberty both with Descartes's hypothetical 'evil genius' which he discussed in the first of his *Meditations*, and Hilary Putnam's (1983, pp. 5-6) notion of the 'brain in a vat.'

out these possibilities, then I could not conclusively claim to genuinely have any knowledge in the way that I had taken for granted throughout my life up to that point.

This skeptical scenario and others like it have been a perpetual thorn-in-the-flesh for theories of knowledge. They gain their momentum by denying both our usual conception of reality (i.e., what we think actually exists independent from us) and our usual conception of knowledge (i.e., our cognitive connection to features of that reality) through the claim that we must be able to rule out every possible alternative explanation for the reality that we think we know in the way in which we think we know it.

Since Descartes, epistemology (the analysis of the concept of knowledge) has been a project heavily dedicated to solving this problem of the conflict between the agonizing pull of skeptical arguments versus our not-too-easily-rejected intuition that do we indeed know things (and a lot of things) because if we didn't we wouldn't be able to get along in the world as well as we do. So, because we do get along pretty well, knowledge is neither an empty nor meaningless concept for us despite our lack of ability in conclusively verifying the non-existence of evil geniuses and envatted intelligence. So, then, how do we reconcile this conflict?

What I think has been one of the most novel and interesting ways of late has been to concede both claims, that is, to allow for the possibility of an evil genius, for instance, but at the same time come to realize that in practical

terms such a possibility is not relevant to the knowledge claims regarding most of our mundane daily life stuff. What the skeptic demands in positing her extreme alternative scenarios are very high epistemic standards that human beings possibly may never be able to meet. This shortcoming, however, shouldn't be taken as an automatic preclusion of the fact that for other kinds of propositions we recognize and employ less intense standards that don't require ruling out skeptical scenarios because such scenarios don't really provide a relevant objection to these kinds of propositions. As Keith DeRose puts it, "the fact that the skeptic can thus install very high standards which we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in ordinary conversations" (1997, p. 917). So, even though there are times when high epistemic standards are the relevant ones given a particular kind of question or knowledge claim (e.g., such as those regarding the nature of our existence, etc.), there are also times in which they are irrelevant (for knowledge claims regarding propositions such as whether or not Tide® is on sale down at the local Piggly Wiggly®). The question, then, becomes this: what is the important factor that makes the difference between instances where high standards are the relevant ones and other instances where they are not?

The answer to this question is that it all has to do with *context*. Epistemologists such as DeRose, as well as Stuart Cohen[2], suggest that we should think about each situation in which we make a knowledge claim as being a particular context whereas each context differs from one another in terms of the particular epistemic standards that must be met in order for a *correct* knowledge attribution to be made within them. Some contexts may require stricter-than-usual epistemic standards due to the fact that they have more pressing and distressing alternatives (e.g., skeptical scenarios, etc.) that are relevant. But in other contexts where such alternatives are not relevant, skeptical scenarios do not need to be painstakingly ruled out, and so the standards in such contexts are a bit lower and more easily met.

For example, consider how different the standards are for attributing knowledge of the claim AB^e ('Andy Ball exists in the way in which he thinks he exists') to myself, depending on the context in which it might be attributed.[3] When I was in the context of that epistemology class I mentioned earlier, it was such that the standards that had to be met in order to correctly

[2] I mention these two because they are probably the most popular proponents of attributor contextualism in the analytic-epistemological tradition.

[3] Of course, this example might be a bit tricky because issues about personal identity are complicated, especially when we take into consideration whether or not identity is something that can be pinpointed to at one particular time-slice or if it is something that exists over a longer stretch of time. I want to avoid these issues and merely talk about what it means to know that I am me (Andy Ball) *in the way in which I think* I exist, i.e. as a body, brain, personality, etc., that is a PhD student in philosophy, and a connoisseur of key lime pie and Kentucky-made moonshine. After all, what Descartes so eloquently shows us is that the possibility of skeptical scenarios does not deny us the certainty of our existence, but they deny us the certainty of knowing we exist *in the way in which we think we exist* since the evil genius could be deceiving us about *that*.

attribute knowledge to my claim AB^e were quite high. Why? Because in that context, possible alternatives such as being a brain in a vat, or being tricked by an evil genius, or being in a dream state (and thus not in 'reality'), are relevant because the whole point of such contexts are to entertain such possibilities for its particular purposes (i.e., settling the question ultimately of whether I exist or not). In this kind of context, then, satisfying the question of whether I know that AB^e is extremely difficult given the *relevance* of skeptical alternatives which gives them license to influence the standards for knowing in those contexts.

However, this same question of whether or not I know AB^e is much easier satisfied in the context, for instance, of being counted 'present' at the polling location where I vote on election day. In that context, the point of whether or not the proposition AB^e obtains is to merely prove that I reside in my particular ward. This question can easily and sufficiently be satisfied by presenting my government-issued identification card which a poll worker will review, and will then checkmark my name on the voter roll to signify that I voted, and will then subsequently lead me into the private booth where I make my selections on my ballot, and then give it to the poll worker to be added to the rest. In this context, I can with extreme confidence be attributed as knowing that AB^e because I easily meet the standards of what counts as

knowing a proposition like AB^e in that context. Evil geniuses and vats are entirely 'non-issues' here that have no bearing whatsoever.

This describes the basic gist of epistemic contextualism. It plausibly allows for knowledge attributions of particular propositions in some contexts (such as at the voting polls) while plausibly denying them in other contexts (such as in epistemology classes), while positing that such *seemingly* contradictory attributions are not really so. It concedes the social fact that in ordinary language we use the concept 'know' in different ways according to the different contexts in which we do epistemic work. It doesn't mean to suggest, however, that such attributions may be arbitrarily made so long as there is a context somewhere 'out there' that provides a plausible arena for them. There are *normative* criteria – call them *standards* or *rules* or *conditions* or *parameters* or *norms*[4] – respective to each context that are the requirements that must be met in order to know a proposition in those contexts. One way to think about how such criteria are determined is as the result of work done by those who attribute knowledge claims to subjects. This is the basic gist of *attributor contextualism*, a leading epistemic contextualist theory.

1.2. Attributor Contextualism

According to Keith DeRose (a leading proponent of attributor contextualism) the standards for knowing in any context are tied to what is at

[4] I will be using these terms synonymously throughout this thesis.

stake in that context. For instance, consider his famous example of how the same knowledge claim can be both true *and* false depending on the context in which it is attributed.

> BANK 1: It's Friday afternoon and I would like to deposit my paycheck today, but as I drive past the bank, I notice that the lines are extremely long, and so I decide to wait until the next morning. My partner says, 'maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon'" (DeRose 1992, p. 913).
>
> BANK 2: The details of BANK 1 are the same except for the fact that I just wrote a very large check for my mortgage payment that will bounce on Monday morning if the funds are not in my account by then, thus "leaving us in a very bad situation." My partner asks if I really do know that the bank is open on Saturday. As DeRose notes, "Well, no. I'd better go in and make sure" (p. 913).

For DeRose, the claim that 'I *know* the bank is open' in BANK 1 and that 'I *do not know* it is open' in BANK 2 are not contradictory. How so? Because "the truth conditions of the sentences of the form 'S knows that p' or 'S does not know that p' vary in certain ways according to the context in which the sentences are uttered" (DeRose 1992, p. 914). The stakes – my purposes and interests – that are involved in my epistemic relation to a particular proposition in a particular context will influence how stringent the standard(s) will be for my claim that 'I know x' to be true in that context.

But what if my claim that 'I know that I exist' (such as AB^e), for instance, is challenged by the skeptic who argues that this cannot be affirmed unless I can *completely* rule out the possibility of being a pawn in a matrix-like world? Must I really meet such a high standard of certainty in order to affirm

my claim? Are the stakes so high that they validate the need for equally high (and seemingly insurmountable) epistemic standards? As DeRose notes, when the skeptic formulates the question this way, she "manipulate[s] various conversational mechanisms that raise the semantic standards for knowledge, and thereby creates a context in which she can truly say that we know nothing or very little" (p. 917). Yet, even though she may invoke such high standards in those contexts where the stakes are high, this "has no tendency to show that we don't satisfy the more relaxed standards that are in place in ordinary conversations" that characterize the ordinary (i.e., low-stakes) contexts we more often find ourselves in (p. 917). As such, the higher-than-usual stakes regarding the questions and claims in an epistemology class may plausibly warrant higher epistemic standards while the lower stakes on Election Day warrant lower standards.

Two important factors are integrally involved in how this particular species of contextualism works.

First, there are *intersubjective* "attributor factors" (p. 919) that "set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true" (p. 921). As such, they are the factors that must be satisfied for an attribution of knowledge (whether to oneself or to another) to be true, and as such, they "affect the truth conditions and the content or meaning of the attribution" (p. 921). They are the epistemic *rules* (or *standards* or *conditions* or *parameters* or *norms*) that have been set for a

particular context and which the putative subject of knowledge must meet in order to correctly be attributed as knowing a proposition in that context.

Second, there are those "features of the putative knower's situation" that DeRose calls the "subject factors" (1992, p. 918). These are simply the facts of a particular situation that "determine whether or not the putative subject lives up to the standards that have been set" by the attributor factors for that situation (p. 922). So, whereas attributor factors are going to set the standards for what make knowledge claims, such as *I know that* AB^e, true in a particular context (i.e., whether the attribution will require meeting stringent epistemic conditions such as what the skeptic calls for or more lax ones), subject factors are simply and merely the facts about the epistemic subject herself and the epistemic situation she is in at a particular time. As such, they are the facts about her epistemic position (what she believes) and maybe facts about the processes she used (or should have used) in a particular situation (hence, maybe what she justifiably believes) as well as facts about the particular background conditions or circumstances in which she exercised those processes and formed her belief(s).

In short, subject factors can be likened to a list of what (epistemically) happened in a particular situation. We might think of them as the kinds of facts that an investigative reporter would want to discover in order to make a judgment about an event at a particular time and place (that is, if investigating whether a case of knowledge happened or not was marketable for the evening

news). Subject factors, then, just tell us whether or not a subject's *I know that x* claim meets the epistemic standards and truth conditions of a particular situation where those standards and conditions have already been set by attributor factors.

What varies with context, then, is what is *meant* by 'know,' that is, how good enough 'good enough' must be in terms of the 'good enough epistemic position' one must be in to *know* in a particular context.[5] So, in BANK 1, where the stakes are relatively low, my epistemic position is good enough to *know* that the bank will be open on Saturday because what it takes to be 'good enough' in that situation is a relatively low standard. But in BANK 2 where the stakes have been raised so much that in order to satisfy the conditions for 'I know,' I must now consider the alternatives that are relevant to my claim and I must be able to rule them out. This incites a much higher standard for

[5] Of course, a question that gets raised here (and one I would like to avoid for the most part) is whether the contextualist is making a claim about the *meaning* of 'know' or rather just about the truth conditions for attributions of 'know.' I tend to think that both are the case, that as attributions of knowledge vary with context (as their criteria are different in each context, some more strict than others, some more lax than others), the meaning of 'know' also changes although it still signifies some sort of *positive epistemic status*. This seems to be Patrick Rysiew's (2001) view when he argues that what is really going on in DeRose's bank cases is that when the subject says that he 'knows' the bank to be open in case 1, he is saying 'know' in a loose sense, and when he says he doesn't know if it's open in case 2, he is saying 'know' in a strict sense (pp. 486-7). Whether the meaning of 'know' varies between a looser or stricter sense, then, depends on what the met epistemic conditions are in each context. An alternative to this view might be Dretske's (1981) argument that "we don't have different senses of the verb 'to know' – a strong sense here, a weak sense there – but one sense with different applications" (p. 368). Dretske's idea, I suppose, is that every application of 'know' is an application of the same concept, but that it is applied in either a stronger or weaker way. If this is the case, it's not clear to me how different this is from Rysiew's argument that we mean 'know' in either a stricter or looser way. The bottom line for both views is that the referents of 'know' are indeed different (in terms of degree) depending on how stringent the standards or conditions are in a particular context.

knowing that the bank is open on Saturday. Here, then, in this context, 'good enough' means something a bit different than it does in BANK 1. It's more stringent. Thus, *in this attributor context*, I do not know that the bank is open.

As Fred Dretske (1981) has argued, "when the stakes go up, the stakes associated with being right about what one purports to know, so does the size of the relevancy set" of alternatives (p. 375). Agreeing, David Lewis notes that "when error would be especially disastrous, few possibilities may be ignored" (1996, p. 556). This is all to say that as the stakes associated with a particular knowledge claim increase, so also does the quantity of plausible relevant alternatives that could defeat that claim. So, in order to possess knowledge in high-stakes situations, the alternatives must be ruled out. For instance, because the needs and interests of the subject in BANK 1 are relatively benign (she merely just needs to deposit her paycheck at some point and doesn't really face any pressing consequences if she doesn't do so anytime soon), the amount of relevant alternatives to her knowledge claim is few. But because the subject in BANK 2 faces the added stress of a mortgage payment that is due very soon (stress caused by the penalties that would result from missing a payment, e.g. late fees, a raise in her interest rate, a negative mark on her credit history, possible foreclosure proceedings, etc.) it is important that this subject be very sure of her claim that she knows that the bank is open on Saturday if she decides not to wait in the long line on Friday.

Being 'very sure' here, then, means being able to rule out the possible defeaters to her claim to know the bank's business hours. Because the stakes are high, the subject needs to be 'on the safe side.' What she should probably do, then, is wait in that long line that evening to ensure her mortgage payment can be made, and hence certainly avoid the undesirable consequences of a missed payment. But if she decides to not play it safe and thus operates from her claim to know the bank's business hours, we are right to question the legitimacy of her claim until she can rule out any possibility that she will have to pay the further price of being mistaken.

What this shows is that epistemic standards vary with respect to our epistemic interests. The attribution 'know' is flexible in the sense that it signifies the attainment of a subject's positive epistemic status in various (and varying) situations where the standards for attaining it are different in each of those situations. This is not to say, however, that knowledge itself is completely relativistic or void of any rules or guidelines. As Heller (1999) notes, what is *relative* in our epistemic endeavors is 'knowledge' not knowledge, and such "relativity comes from the fact that [since] 'knowledge' is vague" the concept of knowledge itself "is referred to vaguely" (p. 117). Thus, what I specifically mean when I ascribe the possession of knowledge to someone "depends very much on linguistic facts" because I will be using 'knowledge' to refer to the knowledge I am claiming she possesses" (p. 117). However, "what is true of knowledge – what is true of the property that I am

in fact talking about in the present context – is independent of any linguistic facts about 'knowledge'" (Heller 1999, p. 117). In other words, even though the term 'knowledge' is flexible enough so that it can be used to refer to a subject's successful epistemic position (whether in a high-stakes epistemically-strict context or a low-stakes epistemically-not-as-strict context), such an attribution will ultimately have no content unless it is a *valid* attribution, that is, unless it meets the underlying epistemic standards that are part and parcel of the particular situation or context in which the attribution is made.

This is simply to say that since each epistemic context is ruled by its particular distinctive epistemic standards, the *difference* in how 'knowledge' may validly refer to one's actually having knowledge will be tied to a particular context's standards as opposed to another context's standards. As such, our knowledge attributions make no sense apart from the context in which we make and apply them.

A promising way to think about this is in terms of our membership in an epistemic community where we attribute knowledge to others and they attribute it to us (and from this, we attribute knowledge to ourselves) based on the various sets of standards that we *share* via our application of them in different contexts. After all, we are communicative creatures who share rules and standards for all sorts of things, from morality to mathematics. Our epistemic practices are no different.

But the more interesting point subtly lurking in the background here is that besides the fact that we're mere sharers of epistemic standards (and, hence, apply them to various epistemic contexts) in order to correctly attribute knowledge in those contexts, it would seem to follow that ***as an epistemic community we collectively determine what the standards for such contexts are***. Just as individuals have a personal *interest* in being epistemically correct in the various high-and-low-stakes epistemic situations of daily life (as the BANK cases show), epistemic communities have a *collective epistemic interest* regarding the various-staked situations that go on in the world. Plausibly, then, the anxiety induced by the demands of those stakes guide our determinations as to what 'a good enough epistemic position' is for a particular situation or, put differently, how high a positive-epistemic-status a subject must have in a particular epistemic context in order to be counted as having knowledge in it.

Stuart Cohen provides a way to think about this in that "the standards in effect in a particular context are determined by the normal reasoning powers of the attributor's social group" (1987, p. 15). This helps answer the question of *how* each context's epistemic standards are determined: by the epistemic community or 'social group' that sets "the standards that operate in [a] context [and] determine which alternatives are relevant" (Cohen 1987, p. 19). For Cohen, such alternatives are 'defeaters' that, when relevant to/in a context, they rout its knowledge claims by exposing deficiencies in the reasons one has

for knowledge claims in that context. These defeaters are neither private nor privileged, but are "*intersubjectively* evident" and "whose relevance is obvious relative to a *socially determined standard*" (1987, p. 8; *emphasis* mine). Even if a subject has "ideally good reasons and subjectively good reasons" she can still "fail to know as a result of failing to meet intersubjective standards" (p. 10). Such standards, again, would lay out how good 'good enough' is regarding one's claim of knowing a particular proposition and thus prescribe how high or *positive* an epistemic status must be in order for that status to count as knowledge.

Cohen further suggests that the "context-sensitive [epistemic] standard" is "a probability measure" understood as "a measure of evidentness" (pp. 20-21). In other words, the standards we establish are determined with a view toward what are the relevant defeaters, and what determines whether or not a defeater is relevant will depend on how probable it is within a particular context. This is why the existence of 7,000 barn façades in a local vicinity is a relevant defeater to the claim that 'I know I am looking at a genuine barn,' whereas the existence of only one façade in a county where there are over 7,000 actual barns may not be as relevant a defeater.[6]

The difference between these two contexts – between one where "the standard of evidence [is] relatively low" (where there are many more real barns than façades) for the claim that 'I know I'm looking at a genuine barn,' from

[6] This is a hearkening back to Alvin Goldman's (1976) barn façade example.

the context where "the standard [is] quite high" for that claim (where there are many more façades than real barns) – is a difference between the probabilities of relevant defeaters in each context (Davis 2004, p. 258). Where the probabilities are high, the standard of evidence required to make the claim true is higher, and where they are low, the evidential standard itself is also low. The probability of an alternative to a particular knowledge claim in a particular context thus may help determine what our standard of evidence will need to be for that context. As such, this plays an important role as a tool in how *we* as knowledge attributors determine what the knowledge conditions and standards must be for a context.

My use of 'we,' of course, is intentional here and is a major point I wish to persuasively make in this chapter, that given the intersubjective nature of epistemic rule determination for a context, it is the epistemic community that determines such rules. As Cohen later points out, for particular contexts "*we* may decide that skeptical alternatives are too remote to count as relevant," but nonetheless it is *we* that determine this (1988, p. 97). Lewis concurs by way of his remark that when we ignore alternatives because they are irrelevant to a particular context, "it is *our* ignorings, not [the subject's] own ignorings, that matter to what we can truly say about [the subject's] knowledge" (1996, p. 561). This harmonizes with DeRose's suggestion that it's the knowledge attributors who determine the epistemic standards for various contexts. But I just want to develop this point a bit further than DeRose does by

understanding *attributors* here in a more explicitly social way as a robust epistemic community.

This, of course, is not to preclude any significance to the individual-apriori considerations at work in how various aspects of the standards are influenced, such as considerations from logic or mathematics (after all, the kind of probability I noted above is arguably an apriori concept that plays an enormous role in context rule-determination). But even if this is the case, there has to be some sort of implicit or explicit agreement as to what the standards are and what tools (such as probability) are useful for determining standards in order to have a 'we' that indeed determines those standards. An *explicit* agreement, for instance, might be something like when a professional scientific community comes together to discuss and determine the professional standards of epistemic conduct that its members will have to follow throughout the course of their research. An *implicit* agreement may be something like what initially takes shape through the actual epistemic practices of a more informal epistemic community, but then through the repetition and bequeathing of that practice to further generations, it becomes the norm.

At this point now, I need to make a crucial distinction explicit that has already been implicit because it will be very important for the rest of this dissertation. I want to argue that there are two contextual aspects of attributor contextualism, and thus, two kinds of contexts: *micro* and *macro* contexts.

1.3. Demarcating the Two Dimensions of Attributor Contextualism

Even though the importance of the community or 'social context' is mentioned in the attributor contextualism literature, more extensive ideas about its role in determining epistemic standards are not really developed. Questions we need to ask, then, should be about how this works and what its implications are. Continuing with a theme from attributor contextualism, we should think about what the stakes are that motivate our knowledge claims, not only in terms of what is at stake *for* individuals (and its consequences for the community), but also in terms of what is at stake for one's community as a whole.

For instance, consider again why the stakes are higher in BANK 2 and then why this instigates a higher epistemic standard for knowledge attributions in that particular situation as opposed to BANK 1. If the subject is wrong, there is a mighty high price to pay: late fees tacked onto next month's mortgage payment, a 'red flag' about the borrower noted by the finance company for future reference, a 'late payment made' notation that will be listed for seven years on the borrower's financial credit report, etc. Simply put, the stakes are high because if the borrower is wrong about her belief (i.e., if she has a false belief yet acts on that belief as if it were true), she will end up in a very bad position. What's interesting, however, is that all of these potential consequences are ultimately social constructs. If enough people miss their mortgage payment in a given month, the bank might lose money and not be

able to pay employees who then cannot shop for groceries and cigarettes. And if enough banks experience this kind of shortfall, we may end up with a 'mortgage crisis' like what has been recently experienced in the States.

Furthermore, what might a history of missed payments say about the subject's character? As a society this interests us as well. This is why we use the services of credit-reporting agencies, for instance. The data they provide allow us to get an accurate sense of how risky a financial agreement might be with a particular person in light of what the stakes are (for both parties) in a particular agreement. This sense of risk, i.e. the stakes involved, then, is what motivates and thus informs the particular terms of the contract. After all, a riskier loan will look different on paper (i.e., its terms and conditions) than a safer one will. But what is more interesting, and the point of this chapter, is that such stakes not only influence the particular content (i.e., the terms) of a contract, but by extension will go all the way to influencing the epistemic standards that will be imposed on the person who will need to form particular beliefs in order to keep and meet its terms, which in this case are the terms regarding making on-time mortgage payments. In other words, the conditions will be such that given the degree of importance of making that payment on time, the need for the debtor's *certainty* in matters surrounding the making of that payment (i.e., whether the bank is open tomorrow, whether there is enough money in one's checking account, etc.) will similarly correspond to a

parallel degree: the higher the stakes, the stricter the epistemic conditions, and vice versa.

My point here is that this is done *socially*. As a society we're able to recognize what the stakes are regarding these kinds of issues and from this we impose the sorts of standards that must be met in order for knowledge to be achieved. In a sense, then, the larger social context, which I am a member of, *sets* the stakes regarding the individual epistemic situations/contexts I find myself in on a daily basis, and as a result, *sets* the epistemic standards and conditions for knowing propositions in those individual situations/contexts.

As Mark Heller notes, such stakes are tied to an epistemic community's collective interests since our knowledge attributions to a particular subject are really at root just our way of affirming that the subject "has the epistemic property we care about in that context" (Heller 1999, p. 118). Contexts, then, are differentiated from one another according to what *our* interests are in terms of the epistemic status we want a subject to have in them: "what *we* really care about is that [the subject] have any one of a set of properties [whether a higher-quality epistemic status or lower-quality ones], any one of which would serve our purposes" (p. 119). So, if a subject truly knows that she exists in the way that she thinks she does both at the polls on Election Day as well as in the epistemology class, this is great! Good for her (and especially for *us* too)! But our interests are such that we don't need her to meet the epistemic conditions we've set for the epistemology-class-context in order to serve our purposes in

the Election Day voting context. To require such high standards would be a sort of epistemic 'over-kill.' Indeed, if enough subjects thought that very high epistemic standards were appropriate in this latter context, our collective epistemic interests wouldn't be served at all.

Epistemic concepts and norms are created with *our* epistemic purposes in mind. Ernest Sosa suggests as much in his claim that "we care about [epistemic] justification because it indicates a state of the subject that is important and of interest to his community" (1988, p. 152).[7] The ultimate reason why the epistemic standards in the epistemology class are so stringent is because *we* recognize that *they need to be so* in order to serve *us* well. It is "the *evaluator's* context – her interests and concerns, what is salient to her, her interaction with conversational partners" that "determines the degree of weight assigned to various aspects" of epistemic situations and hence ultimately determines how 'good enough' of an epistemic status the putative knowing subject must have in order to meet our epistemic needs in a particular situation (Heller 1999, p. 120). To put it succinctly, "the evaluator's concerns play the primary role in the initial fixing of the standards […] the contextual force at work is the evaluator's interests" (1999, pp. 122, 125).

[7] In his example, Sosa noted that such a purpose was "presumably, the state of being a dependable source of information over a certain field in certain circumstances" (1988, p. 152). In chapter two, I will link this early-Sosian argument about the role of macro-contexts with Sosa's current worked out and refined theory of virtue-epistemology, in particular, it's implicit micro-contextualism in the positing of two kinds of knowledge: animal knowledge and reflective knowledge.

Consider an example of how this works practically. Richard Rudner once discussed how socially constructed values influence our most highly prized epistemic practices, even science itself, and how this is especially evident when the stakes are high. Rudner's claim was that since "no scientific hypothesis is ever completely verified," a scientist must then "make the decision that the evidence is sufficiently strong or that the probability is sufficiently high to warrant the acceptance of the hypothesis" (1953, p. 2). But this is going to characteristically require more of a subjective judgment-call on the part of the scientist than an objective calculation.

So, for instance, let's say that a research project is investigating a particular drug that has the undesirable side effect of being potentially lethal in some cases. Of course, the hypothesis that 'the drug is safe to use' cannot be confirmed with 100% certainty. The scientist, then, has to figure out its probability of incurring the undesired side effect against its success rate at treating the illness it will potentially be prescribed for (that is, the probability of its safety versus the probability of its harm) and then make a judgment from that data as to whether or not a confirmation of the hypothesis is acceptable. As Rudner puts it, "obviously our decision regarding the evidence and respecting how strong is 'strong enough,' is going to be a function of the importance, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis" (p. 2). In other words, the *stakes* that would be

involved in 'confirming' the hypothesis will indeed put restrictions on whether or not it can be confirmed.

But here is the more intriguing question: how does the scientist actually make this call? Rudner says she does it by making a "value decision" (p. 3). The epistemic circumstances in which this epistemic performance is practiced are such that the researcher must rely on socially articulated *ethical and epistemic values* in order to properly do her work. Rudner's conclusion, then, is not only that scientists *need* ethical and social values, but also that *they already use them* as a factor in deciding whether or not certain kinds of theories should be accepted, especially those where the risk of error from either confirming or rejecting a hypothesis is just too great.

Now here is my question: how does the scientist know when the risk of error is too great? How does she know how high 'too high' is (regarding the stakes) when the probability of a lethal side effect is too great and thus the hypothesis in question *should not* be confirmed? *Where does she acquire such values in order that they should influence her epistemic practices*?

I would argue that her values are influenced directly by her epistemic community(s) in terms of the values it holds and articulates within that community, which in the Rudner example would include the researcher's scientific community as well as her broader socio-political community. Consider, for instance, that many professional scientific communities have determined what their standards are for the different kinds of research projects

they undertake which are then articulated throughout that community in the form of codified methodological guidelines such as codes of ethical practice, or in the particular way apprentices are trained to undertake the practice, and so on. Members of such communities, then, *directly apply their epistemic community's standards to their work in how they undertake that work.*

What this shows, I would argue, is that attributor contextualism can be considered as a theory that posits ***two kinds of contexts***. Remember that the two kinds of factors noted earlier regarding attributor contextualism were *attributor* and *subject* factors. What the contextualism literature seems to focus on more are the subject factors, that is, the various contexts or 'epistemic situations' that subjects are in when knowledge attributions are made, and thus what it means for the subject to be in a 'good' epistemic position in a particular situation. But what is sometimes ignored (or at least is never made too explicit) is that the *attributor* side of the coin is also a context itself, albeit a much larger one. That there are attributor factors – that is, factors of how stakes are determined and epistemic standards are meted out – implies that there is a much broader kind of context in which epistemic judgments are made that apply to the individual contexts understood as *situational epistemic scenarios*, whether they be an epistemology class or a polling place on Election Day, for instance.

The useful distinction I want to make throughout the rest of this dissertation, then, is between ***macro*-contexts** and ***micro*-contexts**. A *micro*-

context is a particular context/situation in which epistemic performances take place and hence where subject factors are discovered (e.g., contexts such as discussing a bank's business hours with my wife, conducting a research project on a drug's side effects, a philosophy class discussing existence, etc.). A *macro*-context (which can also be thought of as the 'evaluator's context' as Heller had described it) is the broader *social* context which influences and discovers what the stakes are (i.e., what *our* epistemic purposes, needs, and interests are) for each various micro-context and hence determines the appropriate epistemic standards for them. In this way, the macro-context determines what are legitimate knowledge attributions for those micro-contexts.[8] I do not mean, however, to imply that at the macro level a community's needs and purposes are always articulated in an explicit manner. But there does seem to be a sort of epistemic reliance that we share with each other in terms of what we doxastically expect from one another in various micro-contexts.

Attributor contextualism, then, articulates the epistemic reality of our ordinary experience in life, that every usage of 'I know' *does not* mean the same thing; their truth conditions do not all share in needing to meet some universal epistemic standard. Such a standard doesn't exist across-the-board. Yet, experience also tells us that we seem to understand what we mean by 'I know' without having to explicitly articulate all the rules of its various usages because

[8] A somewhat casual way of understanding this distinction as it parallels DeRose, then, is that when DeRose is talking about attributor factors, he's referring to aspects of the macro-context, and when he talks about subject factors, he's talking about aspects of micro-contexts.

each particular micro-context implicitly 'sets the tone' in terms of what it wants. As a result, knowing a proposition truly is a social phenomenon, not only in the sense that we learn things from one another, but also in the sense that *attributions* of knowledge only make sense because of others' normative contribution.

1.4. The Normative Role of the Epistemic Community

As I have articulated attributor contextualism, then, it claims that the rules for knowing vary between various micro-contexts in the sense that how good a position a subject must be in to *know* in any particular micro-context will be relative to that context. Furthermore, those rules are determined by the macro-context. But this raises an important question: how does a community do this in such a way that maintains both objectivity *and* normativity without falling into the abyss of arbitrariness?

Robert Brandom offers some insight here regarding how our social practice of assessing and either endorsing or rejecting each other's doxastic commitments (i.e., beliefs) is a normative practice that sets the standard for what will be expected in future cognitive performances. As Brandom notes, this is an inherently social practice: "there were no commitments before people started treating each other as committed; they are not part of the natural furniture of the world. Rather, *they are social statuses*, instituted by individuals attributing such statues to each other" (1994, p. 161).

For Brandom, to say that a practice is normative, however, is to imply that it is authoritative. The question, then, becomes this: what does it mean for there to be a kind of authority regarding what we cognitively commit ourselves to such that it deems our commitments as being either justified or misguided? Brandom answers this by laying out some important points.

First, our claims "must be understood" as part of the intersubjective game "of giving and asking for reasons" the "fundamental sort of move" of which is "producing a performance that is propositionally contentful that it can be the offering of a reason, and reasons can be demanded for it" (Brandom 1994 p. 141). The idea here is that any claim one makes regarding one of her commitments is a move that is fundamental to, and distinctive of, human rational activity.

Second, in treating another's claim ("performance") as such "is to treat it as the undertaking or acknowledging of a certain kind of [doxastic] commitment" which Brandom calls a "normative status" (Brandom 1994, p. 142). When one displays a normative status "the question of entitlement can arise" (p. 142). There, in asserting a commitment, the subject is also asserting entitlement to that commitment. But just because I may treat a belief of mine as significant and then assert it in order to manifest a 'normative status' to the broader community, does this automatically entitle me to that commitment? Not necessarily. Those of us who are "competent linguistic practitioners *keep*

track of [*our*] own and *each other's* commitments and entitlements" and this makes us 'scorekeepers' of each other's moves in the game (p. 142).

Third, the scorekeeper, then, acts as an epistemic referee. She adopts a "practical attitude" that "consists in the disposition or willingness to impose sanctions" on others' commitments (p. 166). She may "punish those who act in ways they are not (taken to be) entitled to act, and those who do not act in ways they are (taken to be) committed to act" (p. 166). **The key here is that she, the scorekeeper, is the *attributor*,** and her role is the "fundamental" one of this game (p. 166).[9] But what this implies is that when commitments are attributed to others, they are (on some level) accepted or endorsed by the attributor herself, even if this simply means that the attributor merely confirms that the subject has such a commitment. Consider, then, how this perpetually modifies the 'score' in a lifelong doxastic game in terms of "the way in which it changes what commitments and entitlements the practitioners, including the performer, attribute to each other and acquire, acknowledge, or undertake themselves" (p. 166). What it means is that the distinctly human practice of making claims, attributing them, and endorsing them, is both social *and*

[9] Along these lines, Brandom even offers an interesting assessment of reliabilism as a norm in reliabilist epistemology. The ultimate, foundational norm is the community: it is the attributions that we make regarding other's doxastic commitments. For this reason, even if we look to reliability as a norm for justifying beliefs, it is a norm that we have utilized for our purposes in through our practice of scorekeeping. As such, it is not '*the* norm behind our norms' (see pp. 220-221). I will have more to say about this in the next chapter when I take up virtue reliabilism directly.

normative. And furthermore, this, which takes place on the social attributional level, is what *makes* it normative.

Now, let's think about this social element in Brandom's theory in terms of the *macro* context that I have argued is the overarching authority on epistemic matters, at least in terms of setting the norms for what it means to know in each micro-context. How is a *macro* context to be understood as involved in this kind of normative social practice? It would seem that what takes place in this game is a sort of *community assessment* whereas the assessment of commitments and entitlements is *ultimately the social practice of scorekeeping*. But here is where things get tricky. Brandom argues that the community ought not to be (paradigmatically) thought of as a sum of individuals who act as *one*. This is wrong because, for instance, "it is not the community as such that assesses the application of the concept yellow, say, but individual members of that community" (Brandom 1994, p. 38). Rather, our social practice is one that's engaged by *individual* members and yet by *all* members of a community. And in virtue of this particular dynamic it is a communal activity: "assessing, endorsing, and so on are all things we individuals do and attribute to each other, thereby constituting a community, a 'we' " (p. 39).

As Brandom sees it, the primary problem with conceiving of the social context as being personified in one gigantic normative 'we' is that it misleads us into thinking that the 'we' is privileged in such a way that it can never be

wrong. "Objectivity" is indeed "a feature of the structure of discursive intersubjectivity," but this shouldn't be construed as implying that the 'social' is to be contrasted with 'individuals,' and thus, that there is a 'we' versus 'I' dynamic in which the former is *always* right simply in virtue of being the 'macro' context (1994, p. 599). Indeed, "belonging to the community is a concept used so as to have normative consequences of application, concerning the member's being responsible to the assessments of the community, being subject to its authority" (p. 40). But to conceive of and treat the community in terms of these "*I-we* relations rather than *I-thou* relations as the fundamental social structure" is an "orienting mistake" (p. 39).

This fundamental normative social practice – *scorekeeping* – is that of *attributing* to others their commitments and entitlements. This is a practice that, for the most part, is undertaken *between individuals*. Such an *I-thou* relationship, then, "focuses on the relation between the commitments undertaken by a scorekeeper interpreting others and the commitments attributed by the scorekeeper to those others" (1994, p. 599). Understanding intersubjectivity in this fashion allows us to see that the perspective of each individual "is at most locally privileged in that it incorporates a structural distinction between *objectively* correct applications of concepts and applications that are merely *subjectively* taken to be correct" (p. 600).

The question, however, then becomes this: how is it that we separate what is actually correct from what is merely taken to be correct? Brandom says

this is "a messy retail business of assessing the comparative authority of competing evidential and inferential claims" (p. 601). It is something that has to be worked out and negotiated through the practice of *the giving and asking for reasons* and thus becomes what we settle upon and collectively endorse (which is to say, what we as individuals – in the course of our doxastic give-and-take games with other individuals – come to a collective arrangement on).

However, the more important point of Brandom's argument here is that this business is worked out through the social practice of attributing commitments to one another as part of the "implicitly normative practice" of giving and asking for reasons (1994, p. 165). Beliefs get their authority and their correctness, ultimately, from the macro-context – through the distilling and filtering of a sort of general societal epistemic peer review.

1.5. The Epistemic Role of the Normative Community

Brandom's argument, then, gives us a way to think about how an epistemic community, that is, a particular *macro-context*, is a normative one. Through the daily intersubjective exchange of doxastic commitments and entitlements what results is a body of beliefs that has been checked (i.e., scrutinized, inspected, examined) and either accepted or rejected by the broader community. This raises an important question, however. Why does the community care about engaging in such activity? Why does it matter to everyone else what my beliefs are as well as which ones are true or false? Why

does it matter to them that I be *properly* entitled to my beliefs and that my own contribution in the role of a scorekeeper (among other scorekeepers) be fruitful?

Although I will discuss Ernest Sosa in much more detail in chapter two, it is worth noting here his claim that the determination of what the right kinds of conditions are for forming certain kinds of beliefs is something we come to a consensus on and 'mutually rely' on one another for (Sosa 2007, p. 83). Without such mutual reliance, we wouldn't be able to track one another's epistemic successes and thus prescribe the standards or norms for attaining such successes in order to facilitate and encourage future successes. If we didn't do this we would become epistemically stagnant. The *macro*-context, then, plays an important *normative* role in setting the epistemic conditions for the various micro-contexts because the community keeps track of what an epistemic success is and how it can be repeated for members in a community. But here we need to answer the question regarding what a macro-context's epistemic role *specifically* is.

David Henderson argues that the concept of knowledge is a social one and that its purpose is "to certify epistemic agents as good sources for an understanding audience" (2009, pp. 119-120). Knowledge attribution, then, is fundamentally a social practice where we (as a macro-context) verify or endorse others (the subjects of our attribution) as good sources regarding a particular belief in a particular micro-context. In doing this we take on the

social role of epistemic *gate-keepers* who either attribute or deny knowledge to others. This role arises out of "the social epistemic need" that we have for good and "qualified" sources (p. 120). Seemingly, this 'need' for certifying others as knowers is imperative for our own human flourishing. We must be able to discern who is epistemically reliable and unreliable about all sorts of things such as automobile safety engineering, food-crop management, medical care, etc.

For Henderson, what this shows is that "the contextual demands on knowledge […] will be rationally conditioned by the stakes within the communities for which the attributor and interlocutor are keeping gate" (Henderson 2009, p. 126). Who we identify or 'certify' as *knowers* – as reliable sources of information – for particular micro-contexts will depend on the amount as well as the quality of true beliefs they achieve in such contexts. This matters because we "are interested in whether the agent is epistemically positioned so as to render information that is fitting to [our] (community's) interests" (2009, p. 124). Thus, when we 'keep epistemic gate' – that is, when we conduct the practice of evaluating what counts as knowledge (and thus who counts as a certified knower) – we're signifying that *our* stakes are more relevant and more important than any particular individual's stakes. Of course, this assumes that on some level we share the same stakes, that we have the same needs and concerns. But this seems more like a truism than a problem.

We do seem to share some semblance of agreement as to when and why the stakes are higher in some micro-contexts than in others.

Three thoughts come to mind, then, regarding this social aspect.

First, there are many concerns or *stakes* that we all share. We all want the most reliable, certified 'knowers' (i) running the nuclear power plant on the edge of town, (ii) inspecting the food today that we will purchase at the grocery store next week, and (iii) designing the air bags and safety belts that will go into the cars we drive, among other things. So in these particular micro-contexts of nuclear energy, food, and automobile safety, for instance, we work together as an epistemic community to come up with the *requirements* that 'knowers' must meet: they must have extensive training and the proper requisite experience in order to be certified or attributed as such.

Secondly, given the collective stakes, such knowledge cannot be gathered by nor attributed to merely one's self independently of the macro-context. It is something we have to do collectively. As such, we need others to be good at 'doing knowledge.' We need them to go about their doxastic lives in such a way as to be reliably successful at appraising others' reliable successes. And then by certifying such sources, "one is including their claims in a body of accepted results on which others might draw" (Henderson 2009, p. 126).

Thirdly, there would also seem to be stakes that are particular to us as epistemic individuals. After all, each of us has our own life projects and it is plausible that the narrower they are, the more unlikely it is that they are shared

with others. So, it is not as obvious how relevant the 'social' is in such cases, especially since knowledge is a concept that we apply to ourselves quite often. As such, we become reflective thinkers: we learn how to step back and look at our own doxastic practices from the perspective of an attributor, assessing our own evidence or other justifying reasons for our beliefs as well as our reliability at doing it successfully. But the greater point here may be that this is something we initially learn at the *macro* level. We learn it by seeing it done, and then by doing it, maybe even in the way that Brandom has discussed it above. The use and application of *knowledge* thus is ultimately the concept of appraising others' and our own doxastic capabilities that *we* need in order to live well.

Although this discussion doesn't have to do this, it nevertheless reveals that *knowledge* is an evaluative concept. To attribute knowledge is to attribute a success – a job well done – in that it has achieved its goal. It does not make sense to talk about 'false knowledge' (to the chagrin of many of my undergraduate students who think this term is synonymous with 'opinion') as a failure in the way that an archer's shot that misses the bulls-eye (and even the target altogether) is a failure. To attribute knowledge to someone is analogous to saying that the arrow has pierced the bulls-eye. Both are success-oriented performances. Knowledge is a successful doxastic/epistemic performance just as a bulls-eye-shot is a successful archery performance (and a hole-in-one is a successful golf performance, etc.). But it seems plausible to argue that we

wouldn't be able to describe or even understand such successes without having undertaken its exercise first. As Henderson notes, the concept would not have arisen without the associated practice […] the concept arose with a constitutive eye to the demands of successful practice" (2011, p. 85). Knowledge, then, seems to be something we understand by doing it.

So, what is the practice? It is the exercise of gathering the kind of information we need in order to survive and flourish. It would seem that doing this practice, then, would not only give rise to the concept of knowledge, but will also make apparent those qualities one needs to be successful at it: "the character traits, capacities, and skills" that make for one being a good knower (2011, p. 86). Not only will such qualities be identified in successful knowers as those things which lend themselves toward successful knowing, but will then be idealized as those things we *ought* to possess in order to be good sources of information.

And, so, we come back to the role of the community in identifying those knowledge-achieving qualities. *We* collectively and generally agree as to what counts as knowledge for various micro-contexts and thus what are the standards for knowing that a subject must meet, that is, what traits/qualities/abilities/skills she must possess in order to *know* in that context. As Ernest Sosa has noted,

> We are social animals. One's linguistic and conceptual repertoire is heavily influenced by one's society. The society will tend to adopt concepts most useful to it. A concept of epistemic justification that measures the pertinent

> virtues or faculties of the subject relative to the normal for the community will be useful to the community. The community will hence tend to adopt such a concept. (Sosa 1988, pp. 153-4)

One potential issue that arises, however, in thinking about the macro-context as having this epistemic role as '*mutually relied upon*' norm setters (Sosa) or *scorekeepers* (Brandom) or *gate-keepers* (Henderson) is that such a collective will necessarily be reflecting their own beliefs, preferences, and maybe even epistemic prejudices in the macro-contextual work that is brought to bear on various micro-contexts. As such, the question ought to arise as to whether or not *those* beliefs, preferences, and prejudices are themselves the product of further macro-contextually approved norms and standards, that is, whether those agents themselves are also included as targets of their epistemic legislation or whether such legislation is merely the result of arbitrary epistemic enactments.

One way to diffuse this issue may be to just accept the holistic nature of the epistemic practice as presented in Brandom's notion of 'the giving and asking of reasons.' There is a sense in which this practice, as Brandom articulates it, forces our epistemic claims and behavior to come full circle, that is, to subject itself to the norms and standards that it metes out to others. This is not to say that the standards themselves don't ever change. After all, the practices by which knowledge is attained (and the norms governing these practices) change all the time in order to encourage the formation of more *reliable* beliefs, for instance, in order to attain more truth than error. At the

same time, they are not done arbitrarily, but out of a collective agreement motivated by the stakes, needs, and purposes of that collective. As such, there is at the macro level a constant flow of outputs (i.e., norms and standards) and inputs (i.e., feedback) that guide and shape the norms that govern our epistemic practices.

1.6. Conclusion: Knowledge is what the Community is Interested In

What this chapter has attempted to do is to argue for two theses.

First, that attributor contextualism is indeed a kind of dual contextualism. The contexts that DeRose and Cohen focus on are those situations or scenarios that one finds oneself in and how that the standards for each context vary according to how important the stakes are in them. These contexts are what I have termed *micro-contexts*. What I have intended to show, however, is that the 'attributor' side of the coin, that is, the framework that determines what the standards and norms are for attaining knowledge in micro-contexts is itself a context, albeit a much broader one – a *macro-context* – that is a social one. As such, macro-contexts are varied. One such context may indeed include all of humanity, but others will be a bit narrower such as particular religious organizations, professional scientific research communities, educational institutions, and maybe even one's own cultural milieu. From a particular macro perspective, then, each micro-context will be different. Some will be more important, and hence the epistemic stakes will be higher, but

some will be less important and thus allow for a more liberal use of the concept of knowledge.

Second, I have tried to argue for the plausibility of the existence and efficacy of the macro-context itself (i.e. the epistemic community(s) that determine the epistemic standards in the various micro-contexts) by using arguments which show that not only does such a community play a normative role in determining what the epistemic standards are for micro-contexts and that it does this through the practice of a sort of intersubjective doxastic-claim-exchange, but that its primary epistemic role is that of determining what knowledge is in terms of what we need to flourish in life according to our collective epistemic interests and purposes. The upshot of this is that knowledge is a social construct. It is something we learn how to 'do' from the various macro-contexts we are part of and have intellectually matured in.

Chapter Two

Virtue and Context in Reliabilist Epistemology

> *It seems that linguistic and/or epistemic communities conceive of knowledge and, more specifically, justification by reference to community correlated standards. Why is that so? […] We care about justification because it indicates a state of the subject that is important and of interest to his community. […] What sort of state? Presumably, the state of being a dependable source of information over a certain field in certain circumstances. […] A concept of epistemic justification that measures the pertinent virtues or faculties of the subject relative to the normal for the community will be useful to the community. The community will hence tend to adopt such a concept.*
>
> – Ernest Sosa (1991, pp. 275-6)

2.1. Process-Reliabilism's Foundation

Among the vast array of epistemological thought, *virtue epistemology* is a fairly recent contribution that has become quite prominent. It is championed by those who find it to be a better refinement of earlier versions of epistemic externalism as well as a promising alternative to a strict Cartesian-styled internalism. One of its more notable proponents, Ernest Sosa, has been a leading developer of its *reliabilist* variant. His ideas have developed out of the process-reliabilism[1] of the mid-twentieth century that, at the time, had

[1] I will refer to process-reliabilism as a separate and distinct view from virtue-reliabilism. Furthermore, where I simply refer to 'reliabilism,' I am using the term generically and intend

become a popular theory of justification as an alternative to the traditional internalist conception. The traditional conception of justification required the putative knowing agent to possess a conscious grasp of the reasons or evidence (i.e. 'justification') behind a particular belief. On this account, for my belief that 'the table is red' to be justified, for instance, I would need to possess reasons that I could identify and articulate, or at the very least have access to them simply upon reflection. Such reasons, then, would be for *why I have the belief* as well as *what makes that belief true*. This particular conception became known as epistemic *internalism* due to this requirement that the reasons for one's belief must be consciously or 'internally' apprehended by the belief-holding agent (or, at least, could achieve that awareness upon reflection) in order for that belief to be justified. So, if the belief is true, and if that belief is (internally) justified, then upon the standard JTB ('justified-true-belief = knowledge') formula, my belief that 'the table is red' would be a case of knowledge since all three conditions are met.

The advent of process-reliabilism in 1970s, however, challenged this standard conception of justification as being too stringent and then popularized a brand new account of what justification could require. Alvin Goldman, the first prolific advocate of process-reliabilism, agreed with

---

to refer to its common elements both for its process and virtue variants. Note, also, that my generic use of the term 'virtue epistemology' is meant to refer to it as a field or domain of study in epistemology and not as a particular position within it. While Sosa refers to his particular reliabilist position as 'virtue epistemology,' I will refer to it as virtue-reliabilism since there are other positions within the domain of virtue epistemology (namely, virtue-responsibilism) that are not reliabilist in the way that Sosa's is.

internalists that justification is important[2] and there must be some process or property(ies) which confer justification on a belief in order for such a belief to indeed be justified. But Goldman diverged greatly from the common internalist position when he argued that "this does not imply that there must be an argument, or reason, or anything else, 'possessed' at the time of belief by the believer" (1992, p. 106). Rather than clinging to the standard Cartesian inspired "current-time-slice" conception where an agent must possess reasons *right now* in order for her belief to be justified, Goldman proposed a theory where "a belief is justified if and only if it is 'well formed," that is, if "it has an ancestry of reliable and/or continually reliable cognitive processes" (1992, p. 117) where "reliability consists in the tendency of a process to produce beliefs that are true rather than false" (p. 113).

For Goldman, processes such as "standard perceptual processes, remembering, good reasoning, and introspection" pass the test since they are cognitive mechanisms that reliably produce true beliefs for an agent, and so in

[2] This is worth noting given the recent debate on whether justification is indeed a necessary component of knowledge and thus of epistemic discourse. In particular is Alston's (2005) argument to the effect that since there really is no way to objectively identify the feature or property of a belief that epistemologists identify as 'justification' – after all, hardly anyone agrees as to *what* it is yet they all agree *that* it is – then the best thing to do is give up on it. The better alternative is to "focus on a variety of *epistemic desiderata* for beliefs, features of a belief that are desirable from the epistemic point of view, the point of view defined by the basic aims of cognition" (p. 19). While I respect Alston's view and the interesting points he makes, I'm not sure if he even really disposes of justification as he thinks. After all, the desiderata that he advocates – reliability being one of them (of many) – would end up being those kinds of things that lend to belief justification, as it is commonly understood to be. There is an intuitive sense in which it is right to understand a reliable belief-forming process, for instance, as that which creates true beliefs. But when asked whether or not we have any good reasons for thinking that the belief is true, we can reply 'Ah yes, because the process itself is a reliable one at making beliefs that are true.' Hence, the belief is justified.

virtue of their reliability they are "justification-conferring" processes (p. 113). On Goldman's conception, then, "the justificational status of a belief depend[s] on its prior history" regarding processes that explain how that belief came to be (i.e., what created it) and, more importantly, why that belief is true (i.e., what is the high-quality process that didn't just merely create it, but also made it to be true), without depending on the agent's *current* mental or cognitive awareness of these processes (p. 117).

Let's consider further my example belief that 'the table is red.' For the process-reliabilist, the reasons that justify this belief have to do with the reliability of the process that forms this sort of belief: light is reflected off of the table in such a way to reveal its redness, the environmental/air-quality/lighting conditions for viewing objects at that distance is optimal, and then my cognitive mechanism operates in such a way as to take the colored input data and then produce an output belief that 'the table is red.' Yet, while that is all going on, it may be the case that I do not (consciously) possess the reasons that justify that belief such that I could access and articulate them (whereas those reasons are that which simply explains how this reliable visual-belief-forming process of mine generated the belief that 'the table is red'). However, since the belief was well-formed in that it was formed by a reliable belief-forming process, it is justified. Thus, there does exist an explanation that accounts for how this belief came to be. I just don't have access to it right now.

Nor, plausibly, may I ever have access to such an explanation. I just correctly believe that 'the table is red.'

What process reliabilism allows us to say, however, is that I not only correctly believe this, but I *justifiably* do as well. As Goldman claims, "there are many facts about a cognizer to which he lacks 'privileged access,' and I regard the justificational status of his beliefs as one of those things. […] he [can] have a justified belief without knowing that it is justified" (1992, p. 118). If the traditional internalist conception of justification were the correct one, however, then it would seem to be the case that we actually know very little since very little of our beliefs would be justified. But intuitively this seems absurd.

Succinctly put, process-reliabilism purports that a belief is justified so long as there are good reasons (i.e., reasons for why one has the belief as well as why it is true) for that belief, namely that it is formed in such a way that usually obtains truth, *even if the agent isn't aware* of this. This is what Robert Brandom (2000) identifies as the "founding insight" of reliabilism, that "supplying evidence for a claim, offering reasons for it, justifying it, are not the only ways" in which a belief can get justification (2000, p. 98).[3] On Goldman's

[3] The 'founding insight of reliabilism' is "the claim that true beliefs can, and at least in some cases, amount to genuine knowledge even where the justification condition is not met (in the sense that the candidate knower is unable to produce suitable justifications), provided the beliefs resulted from the exercise of capacities that are *reliable* producers in the circumstances in which they were in fact exercised" (Brandom 2000, p. 97).

view, "just as a person can know without knowing that he knows, so can he have justified belief without knowing that it is justified" (1992, p. 118).

Another way to express the same point: a belief can be justified even if the reasons for its justification are *external* to the agent's comprehension *so long as there are indeed good reasons for that belief*. Process-reliabilism, then, as an *externalist* theory of justification, places justificational weight on the reliability of the process that gives rise to beliefs: beliefs that are the products of cognitive processes are justified so long as those processes are themselves reliable, that is, if they create more true beliefs than false ones.

2.2. Process-Reliabilism's Subjectivism Problem

Process-reliabilism, even though an epistemological breakthrough, has not been without its problems. One criticism in particular has taken aim at the seemingly loose and disconnected way reliabilism portrays justification. This criticism asks if it really makes sense to think that a belief can be justified due to a mere 'external relation' holding between the belief and the process that birthed it without saying anything about the agent's own cognitive or epistemic frame of mind. In other words, can one's true belief really obtain justification without any *conscious* guidance or cognitive management exercised on the part of that agent? To hold that such beliefs are justified would seemingly be to imply that beliefs could obtain justification despite being 'irresponsibly' formed and held.

Internalist Laurence Bonjour, in his famous 1985 work *The Structure of Empirical Knowledge*, argued this by way of some interesting thought experiments regarding fictional clairvoyants.[4] Two of those examples suffice for my purposes here.

First, consider Samantha. She possesses the power of clairvoyance and she believes that she indeed has this power even though she doesn't have any reasons or evidence for it. On a particular day, she happens to form (and subsequently maintains) the belief that the President of the United States is in New York City and she holds this belief despite all of the counter-evidence she is exposed to (television news reports, the official White House schedule, etc.) that claims the President is in Washington D.C. on that day. As it turns out, however, all of the counter-evidence was a massive official hoax that was orchestrated in order to prevent an assassination attempt. Thus, in turns out that Samantha's belief was correct due to the reliability of her clairvoyant-belief-forming-process (Bonjour 1985, p. 38).

Second, consider Norman who is also clairvoyant and a reliable one indeed. But where Norman is really lacking is in the fact that that he has neither evidence nor reasons of any kind to support the claim or belief that he is clairvoyant. Nor has he ever possessed any reasons that would 'justify' (in the internalist sense) the beliefs that have arisen from his reliable clairvoyant

[4] Bonjour actually gives seven thought experiments although two are the most relevant for my purposes. See Bonjour 1985, pp. 37-57.

mechanism. Such beliefs, although reliably true, merely 'pop' into his mind seemingly at random. For example, one day he comes to have the true belief that the President of the United States is in New York City. His reliable clairvoyant mechanism is the cause of this belief but he has no articulable reasons whatsoever to support neither this belief's cause nor its truth (Bonjour 1985, p. 41).

Regarding Samantha's case, it is Bonjour's position that there is no way her belief is justified. After all, given the amount of evidence she possesses against her belief that the President is in NYC, she is being "thoroughly irrational and irresponsible in disregarding [that] evidence […] and this irrationality is not somehow canceled by the fact that she happens to be right" (p. 39). For Bonjour, she is lucky that her belief is true. But given the quality of her counterevidence, that belief was never justified. After all, surely the 'process' of believing things in the face of counterevidence isn't a very reliable one, and that is something that the reliabilist would concur with.

But Bonjour also thinks that Norman is similarly irresponsible and irrational in holding his particular belief about the president despite the fact that he wasn't holding this belief in the face of counterevidence as Samantha did. So why is he similarly culpable even though it's not as clear that he should have 'known better'? As Bonjour sees it, "part of one's epistemic duty [is] to *reflect critically upon one's beliefs*, and such critical reflection precludes believing things to which one has, to one's knowledge, no reliable means of epistemic

access" (Bonjour 1985, p. 42, *emphasis mine*). Norman and Samantha are both irresponsible believers since they hold beliefs without having reasons for them. They are *not being critical about their internal doxastic states* and no amount of external reliability changes this. As such, their problem here is that they are lacking a proper subjective (i.e., internal/reflective) relation/connection to the reasons that would guarantee the truth of their beliefs. The connection that they do have to their beliefs' truth is entirely outside of their awareness.

The fundamental problem with process-reliabilism, then, is summed up in Bonjour's rhetorical question regarding Norman: "why should the mere fact that such an external relation obtains mean that Norman's belief is epistemically justified when the relation in question is entirely outside his ken?" (p. 42). The problem, then, is one regarding the proximity between what is purported to make his belief justified and his possession of this 'external' factor. For Bonjour, if some external observer found Norman's beliefs to be reliable because of his clairvoyance, this could indeed be epistemically useful, "a kind of cognitive thermometer" even, but just because *we* have reasons for why Norman's clairvoyance is reliable doesn't excuse the fact that Norman *himself* is not in a position to have, much less offer, such reasons (p. 43).

And so "*from his subjective perspective*, it is an accident that the belief is true" (p. 43). Sure, there may be a reliable cause for Norman's beliefs and they are true in virtue of being produced by that reliable cause, but they cannot be justified unless Norman himself can rule out the fact that from *his* vantage

point these beliefs are merely accidentally true, and thus are a matter of luck. In other words, Norman needs to be able to demonstrate that that which is the ground or basis for his belief's truth is itself what justifies that belief. But he cannot do this. Again, the problem is that there is seemingly no appropriate connection between Norman's cognitive control or perspective and his beliefs' purported justification.

Robert Brandom has also developed this problem of reliabilism's deficiency in accounting for how agents are properly connected to their 'mere reliably' justified beliefs. He claims that one of the temptations of sympathizing with the 'founding insight' of reliabilism (i.e., that the reconceptualization of justification plausibly allows reliable belief-forming processes to confer justification) is to go so far as to hold that "the concept of reliability […] can simply replace the concept of having good reasons for belief – that all the explanatory work […] can be performed as well or better by the former" (2000, p. 100).

But as Brandom sees it, reliabilism itself would be quite incoherent if epistemic agents could not at least be able to "invoke their own reliability" as the justificatory grounds of their beliefs (p. 107). This is because we are all members of a broader epistemic community that engages in the activities of giving and asking for reasons. So, being oblivious to one's own reliability in forming beliefs might imply that we're also oblivious to others' reliability – and then we couldn't ever really learn anything because we couldn't take each other

to be reliable perceivers and believers (p.107). As such, he wants to argue that reliability on its own is insufficient for belief-justification on a *global* level. That is, although we might relegate some kinds of beliefs as plausibly achieving justification by merely being a product of a reliable process, other kinds of beliefs need something more. They need a more proximate connection to the cognitive standpoint and guidance of the agent who possesses the beliefs. In a word, then, they need to be 'subjectively justified' and not merely objectively justified (i.e., by the objective working of a reliable process).

As I previously noted in §1.4, Brandom's view is that knowledge gets its authority or normative power via being the social practice of intersubjectively positing and questioning ours and others' doxastic commitments. Such claims are propositionally contentful, which is to say that they carry justifying reasons, and as a result, such reasons can be (and should be, if we are good at the game) demanded for those claims. When we engage in this practice, we're contributing to the epistemic community as scorekeepers – we're keeping others 'honest' in their epistemic claims (whether by affirming, denying, criticizing, modifying, etc.) and thus contributing to the epistemic welfare of our community. To a certain extent, then, this is going to necessarily entail that beliefs and the reasons behind them will be articulated not only individually, but also more importantly, at the social or 'macro' level. So, to claim that *all* justified beliefs could gain their justification in an

impersonal-to-one's-cognitive-perspective manner, which is what process-reliabilism seems to suggest, seems contrary to what our actual epistemic practice is at the macro level. Furthermore, given our penchant for good reasons and good reasoning, would we really want it to be the case that what justifies *all* our beliefs is a factor that's hidden from our awareness?

For Brandom, then, it seems that 'bare' reliabilist justification should be limited to certain exceptional cases. *All* justified beliefs should not (or practically speaking, could not) be justified on purely externalist grounds. Even though 'reliability' may be a helpful method for justifying certain kinds of beliefs, it *cannot* be the *only* explanation of what justifies the corpus of our beliefs that result from our social epistemological practices (because we indeed practice some sense of epistemic awareness in many of our knowledge claims). But even where reliabilist justification is a relevant and an acceptable substitute in appropriate cases, it is merely just that – a substitute that garners its intelligibility only because an internalist style of justification has *already* defined the playing field as well as the rules of the game. Because of this, mere-reliably-justified beliefs are quite rare. For Brandom, plausible cases of epistemic success "based on reliability without the possession of reasons […] are essentially *fringe phenomena*" (2000, p. 110, *emphasis* mine).

Consider one of Brandom's examples of this 'fringe phenomena' regarding the pottery expert who can tell the difference between Toltec and Aztec potsherds "simply by looking at them" (p. 98). In this example, there are

no features that distinguish one style from the other upon which this expert "can cite in justifying her classifications" (p. 98). It's simply that when she looks at them "she just finds herself believing that some of them are Toltec and others Aztec" and she is reliably correct in her judgments despite that "she does not believe that she is a reliable noninferential reporter" of these potsherds (p. 98). But at the same time, "suppose that her colleagues, having followed her work over the years, have noticed that she is in fact a reliable distinguisher of one sort of pottery from the other" even though she never has fully trusted her off-the-cuff judgments and always would make herself do more thorough chemical and microscopic analysis back at the lab (p. 98). What Brandom claims, then, is to her colleagues "it seems reasonable […] to say, in some cases where she turned out right, that although she insisted on confirmatory evidence for her belief, in fact, she *already knew*" what type of pottery fragments she was looking at "even before bringing her microscope and reagents into play" (p. 98-99). In a case like this, then, "knowledge attributions can be underwritten by a believer's reliability, even when the believer is not in a position to offer reasons for the belief" (p. 99).

In thinking about this particular example, however, it's not clear that it and others like it[5] are truly representative of the kind of reliabilist justification that Brandom suggests it is. He is supposing that this example shows how a

[5] Another of Brandom's examples is that of the chicken-sexer who is extremely reliable at separating male and female chicks soon after they have hatched even though he possesses no justifying reasons as to why he 'knows' how to tell them apart (2000, pp. 102-104).

putative knower can be justified in her belief when she possesses no reasons whatsoever for what it is that makes her belief true. But at the same time, Brandom notes that this expert's colleagues have kept track of her reliability. They constantly remind her of it. They talk to her about her impressive track record at correctly identifying the potsherds and may even tell her that she needn't go to the extent that she does to confirm her knee-jerk beliefs that result from her initial inspection of them. And so, given these aspects of the example, it seems to be a bit of a stretch to claim that this putative knower 'cannot offer reasons' that justify these 'pop into the mind' beliefs because it appears that she is bombarded with reasons. After all, she would seemingly have memory beliefs regarding her own track record (that her thorough inspections of the fragments following her initial judgments 'in the field' have reliably confirmed the accuracy of those judgments time and time again). It also seems that she would possess beliefs regarding what her colleagues have told her about her track record (i.e., their numerous confirmations of her reliability). And hence it would seem that based on these beliefs she would then *have her own* inductive *reasons* regarding her own reliability. So, it's not as though she doesn't possess any reasons whatsoever regarding why she is a reliable perceiver of the fragments. She has plenty (even though she may refuse to cite them, but she's got them nonetheless).[6] As such, it's not clear that this

[6] And this would seemingly go for the chicken-sexer too. He obviously must be aware of his own reliability and track-record at sexing chicks or else he would have lost his job not too

is an accurate case of '*non*-inferential' reliable justification as Brandom suggests that it is.

But what is clear is that, for Brandom, justification can be plausibly reliabilist only when there is a community 'keeping score' (in terms of his 'scorekeeping' methodology discussed in §1.4.) regarding whether or not its members are reliable at forming certain kinds of beliefs that are true despite its possessors' lack of reasons. After all, he notes that it is "unclear" that we could ever make sense of an epistemic community made up of putative knowers who "*never* are in a position to offer reasons for their beliefs. This would require that they never take themselves or one another to be reliable" (p. 107). As such, Brandom appears to suggest that reliabilist justification only makes sense when understood in the broader framework of a sort of social-*internalist* justification of 'keeping each other in check,' so to say, that characterizes the practices of a putative knower's epistemic community. Obviously, this kind of macro work would *necessarily* be somewhat internalist since it requires its participants to be able to converse and debate the reasons behind each others' beliefs. This would necessarily require them, then, to have an awareness of what those reasons are even when those reasons are 'mere reliability.' In this way, then, the kind of reliabilism with which Brandom seems friendly towards

---

long after getting it.

is a tempered reliabilism – an internalist sort of reliabilism – and not a full-blown externalist process reliabilism of a Goldmanian nature.[7]

Nonetheless, despite Brandom's reinterpretation of process reliabilism, the insight he offers in his discussion is quite valuable since it helps explain the role of an epistemic macro-context in its use of reliability as a valid rule or norm for justifying the beliefs of those who themselves have no reasons for their true beliefs at a particular time. Hence, it equips them to able to attribute knowledge to such agents. Understanding this aspect of reliabilist justification as being a sort of 'social' phenomenon, then, gives us an account of one aspect of the sort of macro work that an epistemic community does (and must do) in order to keep track of its members' epistemic statuses.

And, so, even though at heart Brandom's discussion of reliabilism has an internalist tinge to it (given his desire to reduce reliability appraisals to the recognition work of an epistemic community whose members keeps score of each other), this doesn't necessarily preclude genuine cases of the justified true beliefs of certain members of the epistemic community who indeed have *no* reasons for those beliefs (unlike the pottery expert and chicken sexer). Consider, for instance, that very young children indeed know things despite their lack of reasons for those beliefs (and if they're quite young they would

[7] As Brandom explicitly notes his own understanding of reliabilist theories of justification, "assessments of reliability (and hence of knowledge) […] concern the reasons possessed by the *assessor* of knowledge rather than the *subject* of knowledge. […] They should not therefore be seen as external to the game of giving and asking for reason. […] Reliabilism points to the fundamental social or interpersonal articulation of the practices of reason giving and reason assessing within which questions of who has knowledge arise" (2000, p. 120).

lack the linguistic ability that would allow them to even structure a reason cognitively).

An explanation of this phenomenon on Brandom's scheme could be that young children know things because, after all, *we* (their epistemic evaluators and scorekeepers) attribute them with possessing reliably attained knowledge. Consider, for instance, that my eighteen-month-old nephew knows which cupboard is the particular one (among many others) where the Corn Flakes are stored because *we* have taken note of his amazingly reliable (and irritating) track record of finding the flakes only to dump them out of the box and scatter them around the kitchen floor. The ascription of his justification for this belief, then, depends on *others* (i.e., *us*) recognizing that he is indeed justified in his belief when he doesn't even recognize that himself. As such, even though he doesn't possess reasons at the time in which he forms his belief (and hence isn't justified *by his own merit*), his justification is putatively granted by the broader epistemic community that somewhat achieves it for him by citing his reliability and attributing that knowledge to him. And at the same time, this is all part of initiating him into the epistemic community that will continue to do this more and more as he matures cognitively and doxastically until he is finally able to justify his own beliefs at some point.

So, it seems that this example merely reinforces the notion that where there are cases of genuine reliabilist justification, as Brandom argues, their "intelligibility is *parasitic* on that of the reason-giving practices that underwrite

ordinary ascriptions of knowledge" and so "reliability can take a subordinate place alongside reasons […] but it cannot displace giving and asking for reasons from its central place in the understanding of cognitive practice" (p. 110, *emphasis mine*). As such, 'mere reliability' is a plausible explanation for *some* beliefs' justification when it takes place in the larger social arena where the game of 'giving and asking for reasons' is played since such justification *depends* on an epistemic community that recognizes and ascribes it throughout the course of the game. And so, for the putative knower to be *appropriately subjectively justified* via mere reliability really means that she doesn't possess reasons herself (not even of her own reliability). Her justification is 'parasitic' on an epistemic community who possesses those reasons (i.e., her own reliability) for her, and hence serving as her justificatory proxy or surrogate.

So, in this way, Brandom offers an interesting way to think about how putative knowers can be genuinely (and *subjectively*) justified in virtue of their reliability. And, as I will argue in the following sections of this chapter (as well as in the final chapter), there is something very insightful about this notion of an epistemic community that regulates the kind of reliability that is appropriate for determining and ascribing justification in many different micro-contexts.

But I don't think that we have to go where Brandom does in his suggestion that reliabilist justification only obtains where there is an epistemic community of scorekeepers that actually *possess* the reasons that the putative

knower doesn't. After all, it seems that part of the work of the scorekeepers could be that they indeed recognize that some putative knowers (like young children and animals) reliably know things when neither those individuals nor the community at large are aware of the reasons that justify their beliefs. After all, it's not the community itself that *makes* a young child a reliable perceiver (that's what the agent herself does, and sometimes without even being aware of it). Rather, what the community does is *appraise* or *ascribe* her as reliable when it is indeed appropriate to do so according to the norms and rules it has set. So, if this is what the community really does, then it seems that a reasonless knower is still such when the community isn't around to make that assessment for her. So, might there be a way to think about what it would mean to be plausibly justified in one's beliefs upon one's own meritorious reliability even though she (nor anyone else) recognizes this? If so, it would have to satisfy our need here, then, for a way that describes knowledge as being subjectively appropriate from the knower's standpoint or point of view where 'point of view' does not mean 'internalist.'

John Greco offers a comment that points to an interesting solution that may offer us a way of thinking about *merely* reliably justified beliefs that connects the agent subjectively to that reliability. He agrees with the problem proposed by Bonjour and rejoined by Brandom, and notes that "beliefs must be subjectively justified as well as objectively reliable" (2000, p. 97). 'Mere reliability' on the process-reliabilism scheme fails to explain the all-important

subjective-perspective component regarding the putative knower's epistemic standpoint. And, so, Greco asks: "in what sense must knowledge be well formed *from the knower's point of view* as opposed to objectively well formed or *de facto* reliable?" (p. 180).

His answer: that "knowledge has to be subjectively appropriate" means that there must be some sense in which the agent is properly connected to her beliefs in a parallel way to how *a belief's truth needs to be properly subjectively connected to the powers or competences of the agent that justifies that belief* (p. 180; *emphasis* mine). The solution that Greco endorses has us reconsider what reliable belief-forming processes actually are: they're not distant out-of-touch mechanisms, but *internal dispositions* of an epistemic agent. Such an account, then, would give us a way of explaining that for some cases where a putative knower is justified on the basis of her mere reliability, she is justified (despite the lack of being recognized as such) on the basis of *her reliable dispositions* at forming true beliefs. As such, understanding reliable processes in terms of reliable *dispositions of agents themselves* could indeed explain how that children and animals, for instance, could know things despite their lack of being able to fully participate in an epistemic community.

Along these lines is where Ernest Sosa's brand of epistemology facilitates externalism's transition from processes to virtues, and in doing so, provides a theory of justification where a belief can be (1) *justified* on the basis of being formed by a reliable competence possessed by the agent, as well as (2)

'*subjectively appropriate*' to the agent since her belief's justification is a direct result of her own *personal epistemic competence* at forming true beliefs. Sosa offers a more developed theory that answers the question as to how mere reliability itself is plausible for justifying certain kinds of beliefs given our purposes and needs as an epistemic community where such reliability is rooted in an agent's intellectual competences or abilities. As a result, it gives us an interesting way of thinking about what the important factor is that not only justifies beliefs but at the same time is what explains the likelihood of their truth as well.

2.3. From Goldmanian Processes to Sosian Virtues

Process reliabilism, then, offers a way to think about how an agent can have a *justified* belief despite that agent not having any sort of reflective or conscious access to the reasons that justify that belief: it is justified in virtue of the reliability of the process that produced it. As noted above, however, one problem with this is that it possibly makes the justifying factor too distant from the putative knower and thus allows the agent to be in a position where she cannot recognize nor regulate the proper connection between the *truth* of her beliefs and the justifying factor that *makes* them true.

Working from the externalist foundation laid by process reliabilism, Ernest Sosa has sought to solve this explanatory deficiency (or, maybe just its lack of clarity) regarding the *subject's* proper stance that is needed for her belief

justification rather than merely citing the status of her cognitive processes. Sosa's way of attacking this problem results from his thoughts regarding what it means for an *agent* to exercise a belief-forming process or mechanism that forms true beliefs in two kinds of ways: when there is absent any conscious reflection about what justifies one's beliefs as well as when there indeed is such a reflection. And so, while he retains much of the spirit of process reliabilism, Sosa's theory manifests a subtle shift in perspective: instead of thinking of processes as mechanisms that are distant and isolated from an agent's comprehension and thus as something that is separate from the agent (a kind of dualism that a cold and calloused externalist theory might imply), Sosa bids us to think of them as *excellences* or epistemic *virtues* of an agent that, even though she may be unacquainted with their delicate intricacies, are nonetheless part of her cognitive and epistemic makeup. And as such, their achievements – when they successfully work and hence form true beliefs – are to be praised . For Sosa, intellectual virtues are *reliable* epistemic dispositions of the agent and hence are integral in contributing to her competence at successfully forming true beliefs. An exercise of intellectual virtue is an exercise that intentionally seeks to form a true belief. Because exercising such virtues reliably attain true beliefs – i.e., a majority of true ones over false ones – such an exercise, then, confers justification on the beliefs that they produce. So, for Sosa, when such beliefs are indeed true, that which has achieved their *truth* (i.e., the exercise of

reliable epistemic virtue) is at the same time what confers upon them their *justification* (i.e., conferred by the exercise of reliable epistemic virtue).

One of Sosa's initial descriptions of an intellectual virtue is that it's "a competence in virtue of which one would mostly attain the truth and avoid error in a certain field of propositions F, when in certain conditions C" (1991, p. 138). Later, Sosa modified his notion of an intellectual virtue to be understood as an intellectual "skill" that contributes to a belief-performance's success when it is the result of the competent exercise of that skill "in a situation appropriately normal for that exercise" (Sosa 2007, pp. 23, 84). And in his latest work, Sosa refines this even further to understand epistemic virtues as "abilities" that are "a special kind of disposition" of an agent (2011, p. 80). So, whether we would rather define such virtues as competences, skills, faculties, or abilities (or even "capacities" as Alston does),[8] the general idea is that they are some kind of proficiency or aptitude that an agent has, the exercise of which (in conditions appropriate for its exercise) contributes to her epistemic successes (i.e., the true beliefs she creates and sustains).[9]

The relationship between a virtue and a belief, then, is analogous to the relationship between any kind of skill and its successful result (thus mirroring the relationship between a belief-forming process and its output belief). Belief

---

[8] See Alston 2005, p. 154

[9] At one point, however Sosa explicitly notes that he would rather think of a cognitive faculty or virtue, not merely as an 'ability,' but as a 'competence,' insinuating there is a difference between an ability someone might have to do something and the ability to *do it well* (1991, p. 274).

formation is a *performance* that is goal-directed toward achieving truth and justification. However, the value of the performance is not entirely wrapped up in this goal alone, that is, in achieving these epistemic goods independently of each other, but in achieving them *together* as having a *proper connection to one another*. The value of a true belief, then, isn't determined solely because it is true, but in *how* it achieved that truth. For instance, consider Gettier's (1963) example of Smith's belief that the person who would get a job (that he and others had applied for) would have ten coins in his pocket. Why did he form such a belief? Because during the job interview, the hiring manager told Smith that Jones would be getting the job. And prior to the interview, Smith noticed Jones count ten coins and place them in his pocket. But it turns out that Smith ends up with the job. And it also turns out that Smith also had ten coins in his pocket that day too. So, his belief was actually true. But it was true by luck, not by the exercise of any competent process that could properly confer justification on it.[10] So, is it really a highly valuable belief? Did Smith do an admirable, skillful job of forming that belief? Arguably, no. It was an *accident* that it happened to be *true*. The kind of value that we want an epistemic product (i.e., belief) to have, however, is that its truth be the *result* of skillful performance, i.e., that it's success be due to the exercise of an epistemic

[10] Of course, this assumes that luck could not serve as an appropriate kind of justificatory factor, a plausible assumption I think.

competence seated in the agent herself that intentionally seeks to be the explanatory reason why its target belief is true.

Sosa's oft-cited example of the archer is instructive here. Even though an archer's goal may be to strike the center of a bull's-eye with her arrow, the *overall* value of the shot does not simply consist in the arrow making contact with the bull's-eye, but that the archer's accuracy in piercing that very bull's-eye itself result from her *skill* and *competence* at archery. Suppose that a gust of wind were to guide her whimsically released arrow into the center of the target's bull's-eye (Sosa 2007, p. 29). Had the wind not picked up at that very moment, the arrow would have missed the target altogether. But the arrow ended up reaching its target and landing in the very center of it, thanks to the wind. Thus, she wins (had this been a competition).

But what is wrong with this? Is this particular performance *just as valuable* as one where, for instance, the wind might have been blowing in a direction that would have required the archer to exercise some actual skill in 'accounting for windage' in order to successfully complete a shot? Our intuitions seem to say 'no'. No doubt, both performances successfully achieved their target goals – pierced bull's-eyes. But there is something about the latter performance that makes it more valuable. For Sosa, this is because the success

of the latter's performance was due to its being a *competent* performance, i.e. it was a success *because* it was done well (2007, p. 87).[11]

In the same way, then, a belief that is a genuine epistemic success is one that is not merely successful (a true belief) nor merely has been skillfully/competently produced (a belief that's the result of an exercise of epistemic virtue, and hence is justified via this fact), but rather, is a success *because* it is the result of skill and competence: it is "a virtuous performance […] a correct belief *due* to an intellectual virtue" (2007, p. 81; *emphasis* mine). Sosa describes such epistemic successes – i.e., true beliefs whose truth is properly connected to its justificatory (and subjectively so) factors in that it results from virtuous competences – as *apt* beliefs.

For Sosa, an *apt* belief is made up of three important factors. First, it is an *accurate* belief (i.e., a true belief). Second, it is an *adroit* belief (i.e., it is the result of intellectual skill/competence/virtue; as such, it is a justified belief). And third, it is accurate *because* it is adroit (i.e., the belief manifests the proper connection between it's truth and justification; the former is due to the latter) (2007, p. 23). As Sosa argues, "aptness depends on just how the adroitness bears on the accuracy," or, in other words, according to whether or not the skill manifested "sufficiently" influences and contributes to the belief's truth (2007, p. 79). So, when my belief that *p* is true (has reached its goal of being an

[11] "A shot is apt iff the success it attains, its hitting the target, manifests the agent's first order competence, his skillful marksmanship" (Sosa 2011, p. 8).

accurate belief) *because* I successfully and skillfully performed my epistemic skills in forming the belief that *p*, this belief that *p* is then *apt*. For a true belief to be apt, then, requires an appropriate and important connection between its justification and its truth. What kind of connection? In Sosian terms, between its *adroitness* and its *accuracy*; it must be the case that a belief's adroitness *brings about* its accuracy.

A belief, then, can be both accurate (true) and adroit (the justificatory requirement: being competent; manifesting skill), but what distinguishes it as an apt belief over being a mere true belief is *when its accuracy is due to its adroitness*, that is, when its truth-attainment is the result of its justificatory features. This is what Smith's belief (as noted above) was lacking. It was accurate (true), and it was adroit (exemplified some epistemic skill and hence was justified), but its accuracy wasn't due to the particular adroitness that Smith exercised.[12] And thus, his belief wasn't apt.

Although an apt belief's truth is a product of a successful competent performance, for Sosa, this is equivalent to a certain kind of knowledge. He calls it 'animal knowledge,' albeit it is a sort of low-grade object-level knowledge given that the particular competences at play are those that do not require conscious reflection upon their operation. It is shares an affinity, then,

[12] "Something can explain an entity's being in existence without explaining in the slightest why it has a certain property. […] The *existence* of the belief might derive from an exercise of that competence, but this is not enough. It is rather its *correctness* that must manifest the competence" (Sosa 2011, p. 87).

to the kind of beliefs produced by the mere-reliable-processes of process reliabilism (more on this below).

On Sosa's account, however, equally as important as the virtues themselves are the *conditions* under which they're exercised. What happens when the conditions are such that it's not so clear whether or not our virtues could operate properly in them?

Consider the archer again. She is very talented at successfully piercing bull's-eyes with her arrows from 40 yards away. But unbeknownst to her, at a recent competition, a sore competitor tainted her shots with a hidden fan powerful enough to divert the arrows off course every time she took a shot. In this case, the abnormalities of the conditions are to such a high degree that they are not appropriate for the exercise of her particular competence at shooting arrows. As such, there seems to be a sort of circumstantial standard, that is, an understanding of what are the appropriate and normal conditions for the successful manifestation of her archery competence. Under those normal (ideal?) conditions, she is reliably successful. But in these current abnormal conditions, her particular skill-set is not suited to accomplish her particular art successfully. Now, this is not to say that because she wouldn't perform well under these extraordinary conditions she is then totally unskilled or unreliable. As Sosa argues, "failed attempts in abnormal circumstances do not show a lack of ability […] what is required is only that your attempts tend to succeed when circumstances are normal" (2007, pp. 83-4).

Analogously, then, epistemic aptness requires the exercise of epistemic virtue in conditions that are *normal enough* in order to be conducive for their success. As Sosa argues, since "aptness requires the manifestation of a competence," the kind of competence that it must be is the kind "that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it" (2007, p. 29). An epistemic agent's performance of belief formation results in aptness when her competences are exercised in conditions that allow them to work properly. Indeed, when her belief forming capacities are exercised in this manner, she is reliably successful.

Consider my belief that a sofa is at the other end of the room I'm currently in. For this belief to be apt, not only must my vision competences work correctly (such as the cognitive mechanism(s) of my eyes, brain, image processing, and their connection to one another, etc.), but the *conditions* must be normal enough or appropriate for these particular competences to properly operate in the way that they're meant to. For instance, the sofa must be at a reasonable distance in order to view it; there must not be anything blocking my view of it (such as smoke or fog); I indeed need to be viewing an actual sofa (it cannot be a hologram); there should be nothing distorting my perception of it (such as a hallucinogenic drug), and so on. In other words, the conditions must be amenable to my competence's operation: they must be *safe* enough so that exercising this particular epistemic virtue would not repeatedly

result in a false belief a majority of the time. This is explicated in Sosa's condition C:

> For any correct belief that p, the correctness of that belief is attributable to a competence only if it derives from the exercise of that competence in appropriate conditions for its exercise, and that exercise *in those conditions would not then too easily have issued a false belief* (2007, p. 33; *emphasis mine*).

As such, the justificatory factor of an apt belief (i.e., its adroitness), which is the virtuous competence, exercised in forming that belief, depends on the presence of good conditions for belief formation. If the conditions could too easily impede one's epistemic abilities (conditions such as being in a room of holograms, or in a low-visibility situation, or being high on LSD) such that *one's abilities could not be competently exercised under those kinds of conditions*, then one's *true* beliefs under those conditions would have attained their truth by luck. These would be mere lucky epistemic accidents since one's epistemic abilities are not calibrated to such sorts of conditions. To correctly attribute a belief's truth as being the result of epistemic virtue, then, requires that the conditions be conducive for the successful operation of one's epistemic competence or virtue.

So what does this particular epistemic theory get us?

On one hand, we get an account that allows us to attribute a justified belief to Norman contra Bonjour's claim that we cannot. What makes the difference is that Sosa's account could construe clairvoyance as a kind of epistemic virtue or competence for Norman since it reliably gets him true

beliefs as well as properly accounts for why they are true. Because Norman's belief is *virtuously* attained, that is, it is produced through the successful operation of this reliable belief forming *competence* or *excellence* despite his lack of awareness of that competence and its workings, it is apt (and hence, for Sosa, is a case of justified belief that is, albeit, of low-grade quality, i.e. animal knowledge).

But this raises some questions. For instance, since epistemic justification is, on Sosa's account, conceptualized in terms of his notion of *adroitness* – a competent agent-based exercise of intellectual *virtue*(*s*) (i.e., epistemic faculties/abilities/capacities/skills/competences) – must all instances of its exercise *always* guarantee the truth of its output beliefs? If not, then how does this avoid creating more instances of those infamous Gettier cases where one's justified true belief is not knowledge? How does Norman's case avoid being a Gettier case where the truth of his belief is due to luck?

One way to answer this is to note that although a belief may be 'competently formed,' this itself is an evaluative appraisal that will always be dependent on the circumstances under which the competence is exercised. Just like the archer's ability to shoot arrows whether on a calm or on a very windy day, my ability to form beliefs can be exercised virtually under any conditions. But to do so *competently* requires that they be exercised in conditions that at least give my ability a fighting chance at being competent so that my performance may be accurately and adequately appraised. Yet, even when

conditions are normal and my epistemic performance can be adequately appraised, it sometimes may not produce a true belief. So what do we say about those times when my epistemic competence that usually produces true beliefs instead produces a false one? Well, we can say that most of the time, it produces true ones. Blips of reliable virtuous belief-forming faculties are still allowed so long as they are indeed reliable, i.e. they produce true beliefs *most* of the time.

Hence, this helps point out the difference between a justified belief and knowledge on Sosa's account – that is, between a belief that is *adroit* and an adroit belief that is *apt* – in that an apt belief isn't only adroit but it is also *accurate* or true as a result of that adroitness. What Sosa thinks this gets us, then, is an account of knowledge (or apt belief) that avoids Gettier issues because aptness requires this close connection between accuracy and adroitness. It is not enough that a belief be both accurate and adroit (like Smith's belief was, for instance), but it must be accurate *because* it is adroit (i.e., when the adroitness is such that it successfully achieves its intention of forming a *true* belief). For Norman, then, his true belief is not a matter of luck, but is an instance where his clairvoyant ability operates at such a competent level that it produces that true belief. As such, this theory of knowledge is grounded on a theory of belief justification where epistemic virtue answers the question of how subjectively relevant one's belief justification is – one is indeed appropriately connected to her belief's justification because it is the exercise of

her reliable epistemic virtue (i.e., her intellectual/epistemic/doxastic constitution) that justifies it!

Epistemic (or intellectual) virtue, then, is any epistemic competence (or cognitive skill) that indeed is a reliable one: it results in true beliefs most of the time when it is exercised. Such competences may range, for example, from the ability that supermarket doors have to form a sort of belief about when a patron is near and about to enter, to the kind of vision-perception skills that allow a person to form beliefs about what she sees, to correctly completing a logical inference and thus inferring new beliefs, all in a manner that competently forms true beliefs (Sosa 1991, p. 126). So long as such competences influence the formation of beliefs under conditions that allow them to work appropriately and normally, and so long as those beliefs are true *because* they were competently formed, such beliefs are apt. What we find with Norman, then, is that not only is his belief *caused* by his very proficient clairvoyant ability, but also his ability is what *makes* the beliefs it causes *true*. As such, his belief that 'the President is in New York City' is justified (i.e., competently formed from the successful operation of his reliable clairvoyant faculty under conditions appropriate for its exercise) and true, and the latter is due to the former.

What this shows us so far is that Sosa's virtue-reliabilism is thoroughly *externalist*. So long as one's belief-forming mechanism attains true beliefs in a competent (as opposed to a merely lucky) way, such beliefs are apt ones and

thus count as instances of knowledge even if the agent is not consciously aware of the reasons (i.e., her competent epistemic abilities and skills) how and why she came to have those beliefs. Sosa, then, stipulates the term '*animal* knowledge' to refer to this sort of lower-grade or 'first-order' knowledge: it is apt belief that results from the successful exercise of one's 'intellectual virtues' or reliable 'stable dispositions' without requiring any "reflective" or internalist constraint (Sosa 2007, p. 24). As such, stipulating this category of successful epistemic performance allows us to attribute *some* species of knowledge, then, to Norman and others like him.

But this is *not* to say that *all* knowledge is achieved in an externalist way. After all, not all of our belief-forming virtues are directed at forming merely animally-justified beliefs. Some of our competences do us a better service than achieving mere animal justification. Even though some of our epistemic virtues are externalist ones in that they operate absent our awareness of them, others operate in such a way that requires our awareness. These are virtues of a higher order, then. Their successful performance, which involves the skillful exercise of one's conscious apprehension or reflection regarding the reasons for a belief's (i.e., the belief in question) existence as well as the reasons that make that belief true, is the kind of exercise that confers upon that belief a higher-order kind of justification, *reflective justification*. When the beliefs produced by this sort of performance are true as a result of it being a competent one, this achieves what Sosa refers to as '*meta*-aptness' or '*reflective*

knowledge' which is a *better* kind of aptness than mere animal-level aptness due to the fact that it is "better justified" (1991, p. 240).

Why is it better justified? For Sosa, it is of a better quality because the epistemic agent takes a *meta* perspective on her object-level (i.e., first order) belief forming capacities as well as on the conditions under which they operate. As such, the agent has a more comprehensive perspective as to how and why her first order animal belief is true (if it is) as well as to what degree and extent it is adroit, and finally, how the accuracy and adroitness of that belief are connected to one another. Seemingly, then, having this kind of perspective is better than not because the agent is *better subjectively connected* to the particulars of her belief. Of course, cases of mere apt belief (i.e. animal knowledge) do involve enough of such a subjective connection between the agent and what it is that justifies her beliefs (after all, it's *her* competences that achieve *her* beliefs). But with meta-apt beliefs this connection is tighter and more obvious.

Consider Sosa's example of the kaleidoscope perceiver as an illustration of how this works. Subject S visually perceives a red surface and subsequently forms the belief that 'there is a red surface in front of me.' Let's refer to this belief as $p^{\text{RedSurface}}$. Very easily, however, S could have been looking at a white surface illuminated by a red light (and thus would have formed a false belief if this were the case) since there happens to be a jokester J "hiding in the wings" with a red-light-projector at his disposal (Sosa 2007, p. 102). In this particular

case, however, J chooses to *not* manipulate the situation, and so S is indeed looking at a red surface, and so his belief is correct. Sosa wants to argue, then, that S's belief that $p^{\text{RedSurface}}$ is *apt* since it is both accurate (it is a true belief) and adroit (it the product of his reliable visual-perception belief-forming *faculties* working under normal conditions for *its* efficient operation). In light of being adroit in this way, S's belief is justified since the operation of *this virtue* (i.e., perception) under *these particular conditions* (i.e., the conditions are normal/ideal for viewing red surfaces) is a reliable one.

However, S isn't taking a meta-perspective such that would make her aware of this particular epistemic performance and tell us why $p^{\text{RedSurface}}$ is accurate and adroit. S does indeed have the kinds of faculties and virtues that could do this kind of *meta* work (and thus could form *meta* beliefs). But they wouldn't be true. Heck, they wouldn't even be justified. Why? Because the conditions aren't normal for their operation. They wouldn't be reliably successful because the present conditions would preclude this. Although the first-order conditions are normal for S's ability to form vision-beliefs, the meta conditions are such that preclude S from forming any justified true beliefs as to why and how her first order belief is true and justified. Such an account of 'why' and 'how' at the meta level would have to note that S is aware of the jokester who could have made the first order conditions 'bad,' but restrained from doing so. But she isn't aware of this. *If*, however, S was aware of how and why she *aptly* epistemically performed regarding her first-order belief about

the red-surface, then this would be a *meta-apt* performance. Her belief, then, that $p^{\text{RedSurface}}$ would be a better justified and hence a "defensibly apt belief" (2007, p. 24). In the same way and for the same reasons as for this kaleidoscope perceiver, then, Norman's clairvoyant beliefs are apt, but not meta-apt.

Where the major difference lies, then, between animal and reflective justification is between the *kinds* of virtues or competences that are employed for their respective epistemic ends and the kinds of conditions that are necessary for their successful operation. For instance, the achievement of an apt animal belief may be the result of a particular first-order virtue, or combination of virtues, such as physical visual/aural competences (e.g., I see a robin in the tree and also hear it, and because I do so correctly in suitable conditions for doing that sort of exercise, my resulting belief that 'there is a robin the tree' is competently formed and is ultimately successful). However, the ability to reflect on *why* that first order belief is apt is going to require an entirely different set of competences (maybe such as logical inference, coherence, memory) that will determine *how* those first-order competences worked correctly by investigating various aspects regarding how they work in various external/environmental conditions and hence what kinds of conditions would be 'normal' for those first-order competences to operate in.[13]

[13] For the sake of simplicity, noting as examples of first-order processes those competences related to one's physical senses and the kinds of beliefs they contribute to forming, and as

Furthermore, there will have to obtain a set of appropriate conditions for those meta-virtues themselves to operate reliably. Seemingly, then, the successful operation of these reflective-virtues will give us the *reasons* needed for understanding *why* and *how* our first-order beliefs are apt, and thus will confer upon those beliefs a higher order, reflective justification. In short, higher-order reflective adroitness, then, achieves the higher-order reflective kind of justification.

One might point out, however, that this account still allows for some element of luck, especially at the animal level. For instance, isn't the kaleidoscope perceiver 'lucky' that the conditions were such as to allow her vision-belief-forming faculties to work properly? After all, it could have been the case that the jokester had manipulated the conditions, but he didn't. As such, the subject is lucky that the conditions were normal. So, doesn't this Gettierize her belief?

One way to think about an answer to this concern is in terms of the archer again. Consider a particular archery competition she is participating in during a major windstorm. She realizes that she probably won't do her best since her abilities aren't suited to operate well in such conditions. But when she goes to take her first shot, all of a sudden, the wind dies down and becomes perfectly still – and, so, she is lucky that the conditions became

---

examples of second-order processes competences such that involve input beliefs regarding the origin and reliability of one's other beliefs and belief-forming-abilities, may be the easiest way to illustrate the difference between animal and reflective abilities.

normal for her particular degree of ability at shooting arrows. Is this the same kind of luck as what might happen if she takes a horribly inaccurate and unskilled shot that a gust of wind corrects for her? No. Why? Because in the former case, luck merely affects the conditions by making them normal for her to skillfully exercise her arrow-shooting competence *but it doesn't actually affect her competence*. As such, Sosa clams that even though "the act fails to be safely successful since it might too easily have failed through lack of required competence or conditions, it might still be apt, nevertheless" so long as "the conditions remain appropriately normal (or better) along dimensions relevant to the agent's retained competence" (2007, p. 81-2).[14] In the latter case, however, the lucky wind gust actually overrides her own skillful proficiency and is a success *not due to her competence*. The former case parallels that of the kaleidoscope perceiver – S is lucky that the conditions are such that allow her to competently form a belief about the color of the table, but such luck doesn't override or negate the fact that her success is indeed due to her own competence at forming vision-perception beliefs.

On this account of virtue reliabilism, then, one can know propositions in one sense (externalist: 'animal') without knowing in another sense (internalist: 'reflective'), and although each require the exercise of their own

[14] "Our reasoning distinguishes between: (i) factors because of which the circumstances might now easily have failed to be normal, without already being abnormal, and (ii) factors that do already preclude normalcy. Factors of sort (i) make a belief unsafe without precluding its being apt, that is, correct because adroit. Factors of sort (ii) deprive the belief not only of safety but also of aptness" (Sosa 2007, p. 82).

respective kinds of skills as well as their own respective conditions that are necessary for those skills (e.g., the table perceiver will not be able to discern its color if a colored light is obstructing the perception of its actual color, nor will she be able to determine that this object-level deception is occurring if she's not reflectively aware of what's going on with the conditions), both kinds of knowledge share the same basic framework explicated in Sosa's condition C.

2.4. Contextual Elements in Sosa's Early Virtue Perspectivism

While Ernest Sosa has not been totally silent regarding epistemic contextualism, he has also not been explicit in showing how his virtue reliabilism could work under a contextualist framework (or vice versa). Of course, in all fairness to Sosa, this just probably hasn't been an interest of his nor an important issue for his account to sort out. On the few occasions that he has discussed contextualism, he has been careful to distinguish himself and his approach from the more mainstream attributor contextualism. But as I argue in this section, his theory fits much of it very well.

In 1986, as Sosa was beginning to formulate his virtue reliabilist account, he noted that it had the advantage of "conceiv[ing] of knowledge attributions as explicitly or implicitly relativized to an epistemic community (actual or possible) or its corresponding standards" (Sosa 1986, p. 584). He contrasted this "conceptual relativity" account, however, with the "contextual relativity" approach that he attributed to Stuart Cohen at the time (an early

champion of attributor contextualism along with Keith DeRose) whom Sosa noted as arguing for "an explicit or implicit indexical (or some kindred resource), such that the context of attribution determines the community or standards relative to which the attribution has truth value" (p. 584).

So, whereas Sosa himself emphasized attributions of knowledge being *influenced* by community epistemic standards, he understood Cohen to be taking the line that the actual *use* or *practice* of the attribution *determines* (i.e., creates, forms, etc.) those standards. As such, he seems to have viewed Cohen's approach as one that was much stronger regarding the normative role of the context (i.e., *micro*-context). Given this difference between him and Cohen, however, Sosa notes that that he is "not sure how significant this difference may turn out to be" (p. 584).

My own thought on this matter is that Sosa is partially correct here regarding the significance of this difference. On one hand, Sosa and Cohen seem to agree on the importance of a social role in the 'context of attribution.' What Sosa (1986) seems to suggest, however, is that when knowledge attributions are made, they are done so under an already existing community standard(s) for knowledge attainment. So, the standards or rules themselves aren't necessarily 'relativized' whenever an attribution of 'I know that *x*' is made, but rather, they are previously 'relativized' (or conforming) to particular communities in the form of those communities' epistemic standards. This is just to say that different epistemic communities have distinctive epistemic

standards relative to their particular community. What Sosa sees as coming from Cohen's (and probably DeRose's) contextual approach is an argument that the knowledge *attributions themselves determine* the standards: they don't merely verify whether the pre-established standards are met, but they actually create them. So, in other words, Sosa might be assuming that the attributor contextualism championed by Cohen holds that there are no standards or norms governing knowledge attributions until we actually start making knowledge attributions which in and of themselves influence what the standards will be.

While on the surface this may seem like another variation of a chicken/egg argument, there actually may be something more subtly profound going on, which turns out to be a difference in what Sosa and Cohen are thinking about when they think of (and talk about) an epistemic context. For Sosa, the best way to think about the 'relativity' of knowledge attributions is at the social (or what I have termed *macro*) level. (And this will also be a major theme for his 2007 work where he explicitly deals with the role of the community for belief justification). Taking this kind of 'context' approach, then, he is contrasting himself with what he takes to be Cohen's (and DeRose's) '*indexical*' approach, which is what Sosa identifies explicitly as 'contextualism' (i.e., and which actually seems to put the standard-determining emphasis on the role of the *micro-context*).

Of course, as I argued in the previous chapter, I see both of these as two sides of the same coin. They are different aspects of the same contextualist project. So even though Sosa doesn't acknowledge this difference here, but finds the relevant epistemic-standard-creating context stuff happening at the community level (whereas the opposing 'contextualist' view focuses on grassroots attributions themselves), to sharply divide these is unnecessary so long as we have a proper understanding of what they are and how they work. Thus, Sosa was indeed correct in speculating that the difference between the two is probably not that significant.

Two years later, Sosa develops his macro-contextualist line more explicitly. Noting the 'epistemological behaviorism' of Richard Rorty that attempts to explain "rationality and epistemic authority by reference to what society lets us say rather than the latter by the former," Sosa begins to develop an explanation regarding the role that the broader community has in our epistemic endeavors (Sosa 1988, p. 141). As he puts it, "it is believed that something in the context determines the [epistemic] standards" as well as "just how sensitive one must be" to the evidence one has for one's beliefs in order for them to count as knowledge (pp. 143-4). Such rules, he argues, falls on "linguistic and/or epistemic communities [who] conceive of knowledge, and more specifically, justification, by reference to community correlated standards" (p. 152). But the all-important question is *why* this is the case, and

furthermore, how does this link up to epistemic virtue as Sosa conceives of it? Sosa answers:

> All kinds of justification are a matter of the cognitive or intellectual virtue of the subject. *We* care about justification because it indicates a state of the subject that is important and of interest *to his community*. And that holds good for all sorts of epistemic justification, from mere "animal" justification to its more sophisticated, reflective counterpart. In all cases we have a state of interest and importance to an information-sharing social species. What sort of state? Presumably, the state of being a dependable source of information over a certain field in certain circumstances. In order for this information to be obtainable and to be of later use, however, the sort of field F and the sort of circumstance C must be projectable, and must have some minimal objective likelihood of being repeated in the careers of normal members of the epistemic community. For it is through *our* cognizance of such relevant F and C that *we* grasp the relevant faculties whose possession by *us* and *others* makes us dependable informants and cognizers. (Sosa 1988, p. 152; reprinted in 1991, p. 275; *emphasis* mine)

Sosa's argument here, then, seems to implicitly rely on the assumption that epistemology is not merely the playground of philosophers, but that belief-justification and our use of the concept of knowledge itself serves a very practical role in one's community. Explicitly, then, he moves to argue that we as community members are very interested in each other's beliefs being justified. Because of this, then, we're very interested in tracking how reliable we are at forming true beliefs in all sorts of environments and conditions. He thinks this helps us get us at least two things (as he notes in the above quote):

First, the distinction between 'animal' and 'reflective' justification is a useful conceptual tool that helps us categorize and apply different epistemic standards/rules to different kinds of beliefs depending on the kind of belief in question (whether first-order or meta). As such, this equips us to track the

kind of knowledge (and thus, the degree of justification of our beliefs) that we all have and need for particular situations (e.g., Tanis is more reliable at having true meta-beliefs than Dylan is).

Second, that we're concerned about others' dependability means that "we are in no position to require infallibility" but only a "good success ratio" because this evaluative metric will be "most useful to the group" (p. 152). The false need for infallibility (whether perfect reliability and/or a strictly internalist standard) would make knowledge hardly ever attainable and thus useless to us. But when we start with the practical way the community uses and applies epistemic concepts (i.e., when we take a look at the real world), Sosa thinks that what we find is that our 'knowledge' concepts are much broader and that a mere dependability or reliability is all that that we actually require for justification. So, since "a concept of epistemic justification that measures the pertinent virtues or faculties of the subject relative to the normal for the community will be useful to the community," what we'll want to find out is which of these are more *reliable* than others (p. 153). Reliability, then, is what the community is epistemically interested in.

2.5. The Inherent Contextualism in Sosa's Magnum Opus

Now, let's fast-forward to 2007 and to Sosa's most important work that explains the particular brand of virtue reliabilism I outlined in §2.3., *A Virtue Epistemology, Volume One*. There are a few things worth noting from this work

regarding how he invokes the notion of the epistemic community into his epistemological approach.

First, *epistemic standards are community oriented*. Regarding the issue of *intuition* as a kind of apriori knowledge, for instance, Sosa notes that "across socio-economic or cultural divides, we find serious conflicts of intuition" (2007, p. 64). Some cultures value intuition as a way for attaining knowledge, others don't, and the spectrum between these is quite vast. For Sosa, the difference between them may merely be that one side is "saying something about the lack of some relevant communitarian status" while the other "is not denying that, but is simply focusing on a different status, one thought desirable " (p. 65). Here, then, is where a particular (macro) context matters. Depending on which epistemic group one is a member of, the value of various epistemic statuses will depend on what the 'communitarian' standards are, which is to say that it will be based on what is valuable for the community. Why? Because one's epistemic community already is concerned with all kinds of epistemic values. For instance, as Sosa argues, justification is the kind of epistemic good that concerns values such as "being true, being a truth tracker, […] being virtuously based through a virtue recognized as such in the believer's community (and, perhaps, *properly* recognized as such)" and so on (pp. 68-69). In this way, then, the macro-context identifies values and in doing so shapes what the epistemic goods are for the community as well as what the requisite standards for achieving them are.

Second, the *degree of aptness* of one's beliefs will depend on macro-contextual factors. We have seen that on Sosa's approach, for a belief to be apt the epistemic virtuous disposition that produces the belief must also be that which makes the belief true, i.e., it must be accurate, adroit, and accurate *because* adroit, *not* accurate because of luck. But aptness is not an all or nothing affair. It seemingly can be construed (and even exercised) in degrees. What determines the degree of aptness, then, "depends on just how the adroitness bears on the accuracy" (2007, p. 79). For instance, it may be plausible that an extremely long shot that an archer takes, even though it gets *some* help from the wind, might still be apt, but not as apt as it would be if it had been successful without any help from the wind at all (p. 79). So there is a sense in which some performances might be less apt, some more apt, all depending on how sufficiently adroit the agent's skill and competence contributed to her performance. But what this requires is an "index of sufficiency" that contains "some threshold" that allows us to determine how sufficient one's performance is and whether or not her success is completely or partially attributable to her own competence.

So how do we figure this out? Sosa argues this would probably have to be "contextually determined" (p. 79). So, the epistemic (macro) community will determine the appropriate threshold level at the 'low end,' that is, determine when one has acted sufficiently to merit the attribution of a 'baseline' adroit performance. Additionally, the community will figure out

what the threshold at the 'high end' must be for determining when one has performed in a more fully sufficient or *meta*-apt way.

Third (and related to the previous point), those aspects of the conditions or circumstances that might *contribute to* one's epistemic performance by actually helping a competence in some sense (e.g., the slight guiding wind that doesn't, of course, override the archer's competence) in a given situation may not automatically disqualify an agent's performance, but are taken into account for the purpose of evaluating and determining just how reliable the performer is under different kinds of conditions/circumstances. As Sosa puts it, "a performance can still be apt when its safety depends on a circumstantial contingency, provided it is guided by that contingency" (2007, p. 83). In other words, one's performance doesn't need to be completely 'safe.' One needn't always know that the conditions are perfect for one's performance before one actually performs. For instance, "even if the light alternated quickly and randomly between being good and being bad, one can still acquire perceptual knowledge so long as the deliverances of one's color vision are accepted, not at face value, but guided by the ringing of a bell […] only when the light is good" (p. 83). So even if some *contingency* contributes to my epistemic competence in some way, whether it be a ringing bell or a guiding wind, it doesn't take anything away from that competence per se. However, my epistemic community might consider such a belief to be less apt, but probably not non-apt.

The point here is that such contingencies might be acceptable but *only because they provide a useful service to the epistemic community in helping it realize its epistemic goals*: it is in its best interest to know when individuals are epistemically adroit or not (and reliably so), as well as when and where their epistemic competences are weaker or stronger, better or worse, etc. For Sosa, "our purposes in evaluating people plausibly help determine that distinction, given our need for coordination and mutual reliance, and hence for keeping track of strengths and weaknesses, our own and others" (2007, p. 83). The 'distinction' that Sosa is referring to here is the distinction between attributing others (and even ourselves) as epistemically reliable or successful under certain circumstances versus unreliable and unsuccessful under others. What the appropriate (as well as inappropriate) conditions/circumstances are, relative to our epistemic competences or abilities, is something *we* discover as a community. This is why "we value as 'aptitudes' certain abilities relative to certain background conditions" where such aptitudes "are relative to distinctive correlated parameters at the time of their exercise" (p. 83).

So, we might agree that certain kinds of epistemic virtues would be very unlikely to result in apt beliefs if exercised under certain conditions, such as our vision-perception competence of identifying exotic birds on a very foggy day. But given conditions such as foggy days, we really don't expect someone to get the belief right with a high degree of accuracy: "failed attempts in abnormal circumstances do not show lack of the ability. […] What is required

is only that your attempts tend to succeed when circumstances are normal" (pp. 83-4). And, so, when such attempts do succeed, as a community, this is when we note that an epistemic competence or virtue has been successfully exercised under particular conditions, and thus is possessed by a particular agent to the extent that she is successful in performing it under those conditions.

What we do, then, as an epistemic community is come to an agreement as to what the appropriate conditions are for our epistemic competences to work aptly, that is, to work in such a way that reliably produces true beliefs that are true *because* those competences worked correctly. But this will be determined by what our needs are as an epistemic community.

And, of course, these needs may vary. For example, we seem to have determined that regularly being in a 'drunken stupor' is not a normal condition for driving a motor vehicle and would actually preclude one from competently and successfully driving that vehicle safely (even though, however, in places like Toronto such behavior doesn't preclude one from being mayor). Yet, this doesn't preclude that same person from competently performing this task when conditions are normal, such as when he has sobered up.

But the community doesn't just stop at linking up our epistemic competences or virtues and the appropriate conditions for their exercise. Given Sosa's more nuanced and sophisticated epistemological theory, the community does this at different levels. For some micro-contexts, then, the community

might simply require mere apt belief (animal knowledge), but in other contexts, a bit more aptness is called for. For instance, whereas when driving a vehicle it is 'good enough' to have animal knowledge regarding how the brakes work (after all does one really need to be meta-justified regarding why they worked appropriately?), while in a religious or public-policy context, we (hopefully) want articulated justifiable reasons to back up whatever knowledge claims are being made. Why such a divergence? Why do we want something more in some cases and something much less in other cases? The answer, seemingly, is the same one we encountered in the previous chapter: it all depends what our interests are in terms of what's at stake if we're wrong versus what's at stake if we're right.

Fourthly, then, Sosa's approach to epistemology is solidly contextual in that it is both *macro-contextualist* and *micro-contextualist*.

Sosa's virtue reliabilism is *micro*-contextualist since knowledge means different things depending on the kind of epistemic performance successfully accomplished by the subject. A lower-level successful performance will be one that is merely of one's particular belief forming mechanisms working 'under her internal radar,' so to say, that competently produces a true belief and achieves knowledge of the animal sort. But if such a mechanism is indeed 'on one's radar,' which is to say that the agent is aware of it and of the conditions under which it has competently been exercised, then its resulting true belief qualifies as the higher sort of knowledge – reflective knowledge.

Sosa's proposal, then, is that there are different *kinds* of knowledge that are respectively achieved under different *kinds* of conditions (albeit the overall framework – condition C – is similar for each kind). Furthermore, who the subject is, and in what circumstances the subject may be in, will also affect what the standards (i.e., the rules that will have to be met) will look like for the different kinds of knowledge in different kinds of micro-contexts. For instance, the kind of epistemic performance *we* require of a brain surgeon when braking her car is going to be a much different one from the performance we want her to exercise when she is removing a glioblastoma from a patient in an operating room. But at the same time, both of these performances will require different epistemic-criteria-satisfaction than what we'd require of automatic supermarket doors (that is, if it is plausible to attribute such doors as 'knowing' when patrons are about to enter). At this level then, the contextual features of Sosa's account focuses on subjects.

Sosa's theory is contextualist, however, not only in proposing that the attribution of 'knowledge' can be relativized as either stronger or weaker depending on the *kind* of performance undertaken by the subject, but it is also contextualist in that the *parameters* or *rules* themselves are determined by the epistemic community, that is, by the attributors. As such, it is *macro*-contextualist. The epistemic community *sets the criteria* for what is required in the various micro-contexts in terms of the kind of attributions we need to make about the epistemic status, for instance, of the brain surgeon as opposed

to the supermarket door. This would also get cashed out in terms of whether animal knowledge is good enough in some contexts, or if reflective knowledge is what we require.

So, for example, we mutually agree that '*x*' (vision-perception) is an epistemic virtue or aptitude to form true beliefs under conditions '*y*' (clear day) and *not* under conditions '*z*' (foggy day) unless the circumstances admit of some variable that indeed allows the agent to competently see despite the fog (e.g., it just so happens to lift for the five seconds it takes the agent to form the true visual-perception belief). Why? Because under *y* and not under *z* will *x*-beliefs be reliable, and, as noted earlier, reliability is very important to us. It is by gauging others' (as well as our own) reliability in their epistemic attempts that we figure out what sorts of conditions are more suitable for certain kinds of epistemic performances and thus form the epistemic criteria for different contexts accordingly. The purpose of reliability isn't just to 'externally' justify beliefs (as with the old process reliabilism), but ends up being a standard that we employ to evaluate others epistemic performances and thus 'track' *our* successes and failures.[15]

[15] One interesting difference between virtue reliabilism and process reliabilism, then, is that whereas the latter focuses on the reliability of processes/mechanisms of an epistemic agent, the former seems to consider the reliability (or unreliability) of the agent herself and then will look at the process she utilized in forming a certain belief – the *competence* employed under a particular set of *conditions* or circumstances. I'm not sure, however, if this distinction really points out anything significant or is just a different way of expressing the same idea.

2.6. Sosa's Explicit Critique of Epistemic Contextualism

However, despite my argument that Sosa's virtue reliabilism is inherently contextualist, his latest work explicitly conveys his worries about contextualism in epistemology. We somewhat noticed this in Sosa's 1986 work which I discussed earlier, where he contrasts his approach from Stuart Cohen's 'indexical' approach (see my § 2.4. above). There, I noted that Sosa's concern is with Cohen's claim that (*micro*) contexts themselves determine or shape epistemic standards. However, Sosa himself seems more inclined to argue that the community (*macro-context*) is what shapes or determines the standards that get played out in various situations (*micro-contexts*).

Now, fast-forwarding to his 2011 work, Sosa expresses a concern with an epistemic contextualism that is taken too far in order to justify attributions of knowledge without enough scrutinizing as to whether or not knowledge actually has been achieved. As he lays out his concerns, Sosa criticizes contextualism *if* contextualism implies a complete and total relativity of the use and application of the concept of knowledge. Of course, as argued above (and in the previous chapter), such a robust relativity cannot be the case due to the all important macro aspect of epistemic contextualism: the community serves as the arbiter of epistemic rules and thus is its standard-bearer and enforcer. So why might he be concerned about a possible runaway micro-contextualism?

As Sosa describes it, epistemic contextualism (EC) is the thesis that "sentences of the form 'at *t*, S knows that *p*,' are truth-evaluable only relative to

a context of use" (2011, p. 98). In other words, determining whether a subject knows a particular proposition will itself be evaluable in light of the required standards for a particular context (chapter one review: this is why the subject is attributed knowledge about the bank being open in a low-stakes context but denied such knowledge in a high-stakes one). But Sosa argues that EC is susceptible of 'overreaching' here and will inevitably do so when it's guilty of committing the *contextualist fallacy*, which he describes as "the fallacious inference of an answer to a question about the correct use of words in its formulation" (2011, p. 97). EC, then, is culpable of creating a confusion surrounding the notion of what it *means* to 'know' something.

How does this cash out? Sosa seems to think that some contextualists erroneously take it as given that "even if we fail to know about ordinary matters in philosophical contexts, such as whether one has hands, we do often enough know those same matters in ordinary contexts" (2011, p. 104). In other words, even if we don't know some proposition in a high-stakes context, for instance, there's got to be some context out there somewhere in which we do. Even though this is sometimes (or maybe most times) the case, to infer that it *must* be the case is to infer fallaciously. Whether we know such things in ordinary contexts can only be determined once it has been investigated *as to whether or not we actually satisfy the epistemic rules/norms that are set for those contexts*, which of course is going to require the subject to believe aptly according to those rules (i.e., her belief is not only going to have to be true, but

true due to a competence under certain conditions where the 'threshold of sufficient adroitness' will be macro-contextually determined).

So, is there a place for EC? For Sosa, "contextualism gains epistemic relevance if the pertinent contextual variation *concerns only the required measure* of a certain *shared* desideratum" (2011, p. 101 *emphasis mine*). In other words, contextual relativity is applicable only if what is relative or flexible is the measure or amount of a certain agreed upon epistemic 'ingredient,' so to say. What I think Sosa is getting at here are two important points.

First, even though the particular norms will differ between each micro-context, the same sort of rule-governing structure for assessing and attributing knowledge claims should be uniformly applied to each instance of knowledge attribution in those various micro-contexts. This ensures that various normative epistemic criteria and concepts are shared *among contexts*. So, it should never be the case that in one particular micro-context, a knowledge attribution and/or assessment of a knowledge claim doesn't consider whether a belief is adroit or accurate, for instance, when other micro-contexts do consider these particular epistemic goods. Looking for these "various dimensions" in the beliefs that we assess in different micro-contexts should reveal a difference between them *only* in the sense that "each admits a threshold" that is appropriately respective to the context in question (p. 102). The relevant difference between micro-contexts, then, is to what extent these criteria apply. In other words, what 'shifts' between micro-contexts are not so

much the meanings or even the requirement of having these normative concepts, but the degree to which they apply to a particular belief in question.

So, EC is not a thesis about whether or not we should be concerned about the epistemic 'dimension' of justification-in-terms-of-reliability, for instance, but about various micro-contexts, that even though they may differ in terms of how reliable a belief needs to be in order to be justified, they nonetheless are uniform in applying these sorts of epistemic criteria "along the same dimension" (p. 103). So, the meaning of an epistemic desideratum, such as reliability, can *mean* the same thing between different micro-contexts, but such contexts then become distinctive and different from one another in terms of the degree of reliability they respectively require in order to achieve justification in them. The happy upshot of this is that "even if our beliefs do not attain desired levels along certain dimensions, they may attain lower levels," such as Sosa's category of animal knowledge or mere apt belief (p. 103).

But, again, *this is not a given*. To assume that because the degree of reliability of S's perceptual-belief-forming-process required for S's belief that *p* to be justified doesn't meet the ideal threshold required for the skeptic automatically means that it does in some ordinary context, is to attribute knowledge to S *without considering whether S has indeed met the threshold requirement in that particular ordinary micro-context*. And this is fallacious. Furthermore, what this may imply is that epistemic norms and rules in some

contexts might be too low and hence inadequately reflect our shared normative epistemic needs and standards. As Sosa puts it, "from much discussion with undergraduates and ordinary folk, I am convinced that the term 'know' and its cognates are sometimes used as to make it true that the medievals just 'knew' that the earth was flat" (2011, p. 105). I've had similar experiences with students from my introductory philosophy courses who want to say that some things they 'know' are false, while other things they 'know' are true. They have no problem with using terms like 'false knowledge' and 'true knowledge' (i.e., they often conflate the concept of 'knowledge' with the concept of a 'belief').

Sosa's point is that just because knowledge is *attributed* in a micro context doesn't automatically mean that it has actually been *attained*: "that some sophomores call it 'knowledge' hardly suffices to make it so, even if the attribution is correct in their context, by their definition" (p. 105). After all, one might (hopefully) argue that a belief cannot be apt unless the belief is true. So, just because a particular micro context might not require truth for a belief to be apt doesn't mean that that context gets some kind of exemption from 'truth.' Similarly, just because lower thresholds might be the more relevant metric for our ordinary epistemic concerns than philosophical high-thresholds, this "is not something that goes without saying" (p. 106). Rather, "it has to be considered, and perhaps argued, case by case" (p. 106). In other words, it has to be something worked out in (and by) the macro-context. And this leads to the next point.

Second, it is something *we* share. As Sosa notes, the most interesting questions raised by EC are "what are the most appropriate dimensions along which a belief must be assessed in determining whether it qualifies as knowledge?" and "what are the dimensions that *we care about* when *we want* our beliefs to give us knowledge?" (2011, p. 103 *emphasis mine*). This hearkens back to Sosa's explicit argument on the role of the community in setting necessary threshold levels for reliability (and hence, justification) that are required for particular micro-contexts. And further, this is what makes the lower 'ordinary' thresholds acceptable and desirable to us, because as members of an epistemic community, we have an interest in being able to use our shared concept of knowledge in such a way that works well across the spectrum:

> Working from the same nexus of concerns, we wish for beliefs that are at least somewhat well justified and somewhat safe and somewhat assured, and we prefer such beliefs to those that fall below them in those respects, and we prefer this even in cases where we fall short of wished-for heights of assuredness, safety, and rational justification. If so, then the fact that 'knowledge' is correctly applicable in line with the lower ordinary thresholds is indeed relevant to the nexus of concerns that includes our desire for the epistemic heights (Sosa 2011, p. 107).

2.7. Concluding Remarks

What we find in a Sosian-styled virtue reliabilism (as I interpret it, anyway) is a blatant contextualism. Interestingly, then, in his virtue-styled revision of the older process reliabilism that presents an externalist theory of justification that better handles the need for agents to be *subjectively* justified in their beliefs, what Sosa offers us is an analysis of knowledge that puts a

primacy on context: different micro-contexts call for different kinds of epistemic competences in order to successfully 'know' in those contexts. However, we can only make sense of this in virtue of being a member of an epistemic community that highly values epistemic goods such as reliability and truth. Because we're interested in our beliefs being consistently correct, we're interested in others being reliable knowers, that is, reliable in exercising their epistemic competences in all sorts of micro-contexts, whether lower-order object ones or higher-order reflective ones.

Chapter Three

Virtue and Context in Ethics

> *Now, that we must act according to the right rule is a common principle and must be assumed – it will be discussed later, i.e. both what the right rule is and how it is related to the other virtues. But this must be agreed upon beforehand, that the whole account of matters of conduct must be given in outline and not precisely, as we said at the very beginning that the accounts we demand must be in accord with the subject natter; matters concerned with conduct and questions of what is good for us have no fixity, any more than matters of health. The general account being of this nature, the account of particular cases is yet more lacking in exactness; for they do not fall under any art or precept but the agents themselves must in each case consider what is appropriate to the occasion, as happens also in the art of medicine or of navigation.*
>
> – Aristotle (*NE*1103b30-1104a10)

3.1. Taking Stock

Now, we turn from epistemology to ethics. My main concern in this dissertation is epistemology. However, because I'm interested in the responsibilist variant of virtue epistemology, and because it is intimately connected to virtue ethics, I will discuss virtue ethics at some length in this chapter order to motivate the epistemological claims I want to make in later chapters about virtue responsibilism. There are many different ideas about the nature of the connection, but for the most part virtue responsibilism suggests

that intellectual (epistemic) virtues are akin to moral virtues in that they are character traits or dispositions of persons rather than mere abilities or competences at performing particular kinds of cognitive tasks. This is how Linda Zagzebski (1996) frames her epistemic responsibilism and hers is instructive because it is probably the most exhaustive one to have been articulated in recent years. Because of this, her theory is the paradigm of virtue responsibilism that I want to use for my purposes in the following chapter.

Because Zagzebski's epistemic theory is patterned loosely after Aristotle's moral theory, I want to turn to Aristotle first, and specifically, to his *Nichomachean Ethics*. If I can plausibly interpret parts of Aristotle's theory (and what I take to be some of its implications) in a contextual way (as I have defined 'contextual' in chapter one), then Zagzebski's responsibilism may also similarly be interpreted in a contextual way.

There are other reasons, too, why I want to consider Aristotle.

For one, because, after all, it's Aristotle. His is arguably the earliest best-worked-out virtue theory and this is revealed in the fact that it has influenced so much of the way in which all post-Aristotelian virtue theories (to some degree) take their shape. I say 'best' worked out and not 'exhaustively' worked out because *by his own admission* Aristotle's theory is not a complete one. One might even argue that it is something of an ethics without much content since it offers a framework for making moral attributions without offering any precisely delineated decision procedures for making moral

attributions. However, this supposed incompleteness, I think, is key to proposing a sort of contextual aretaic view.

As such, secondly, I want to take some of Aristotle's ideas as jumping off points for a sort of neo-Aristotelian position that has a place for the sort of macro/micro-contextual distinction that I proposed in chapter one while at the same time being charitable to Aristotle in understanding that he is not necessarily a contextualist himself regarding virtue.[1] But what I take from his ethical framework is the notion that there are fundamental moral principles that have to be settled upon and collectively agreed to as starting points (i.e., which I would argue is the role of macro work) as well as various specific manifestations of virtue that will *look* very different depending on the situation (i.e., micro contexts) that call for a virtue, such as generosity or courage.

After suggesting my contextual re-interpretation of Aristotle's basic ethical structure, I will then consider aspects of the virtue theory of Alasdair MacIntyre in order to show some further neo-Aristotelian affinity to my thesis here. MacIntyre's theory will serve as an example of a respected neo-Aristotelian theory that indeed also preserves some of the spirit of Aristotle's ethics without buying into his entire program. Where MacIntyre's thoughts help my overall argument, however, is by what I think to be his more explicit

[1] After all, he is quite stringent about virtue. Either you have it or you don't. And, furthermore, either you have them (all) or you have none. Virtue is an all or nothing affair. It's not as if one can merely pick and choose which individual virtue they want to exercise. If one *is* virtuous (and hence, they possess *phronësis*), then she will *act* virtuous no matter the situation and no matter what particular virtue may be called for in that situation since "with the presence of the one quality, *phronësis*, will be given all the virtues" (1145a1-3).

(yet not entirely explicit) employment of a *macro*-contextual and *micro*-contextual distinction (and connection).

My point in this chapter is to suggest that, at some basic level, an *aretaic* framework can plausibly be embedded with the notion that the broader social context ascertains virtue, both in *defining* it and in *determining* its standards in how it applies to the various moral situations we find ourselves in. To support this, however, we do have examples of how a macro-context might plausibly be understood in terms of moral exemplars (Aristotle) as well as in terms of intersubjective moral communities (MacIntyre).

3.2. Starting with Aristotle

One of the more interesting facets of Aristotle's moral theory (articulated in the *Nichomachean Ethics*) is that it is not merely a treatise of how to discern the morally right-thing-to-do (such as what seems to be the point of much of contemporary ethics), but it is one aspect of a larger, more important project for him. It is what lays the groundwork for a robust political science. His discussion of how to best manage and rule one's *polis* (which is more thoroughly dealt with in his *Politics* following the *Nichomachean Ethics*) really begins in the *Nichomachean Ethics* with Aristotle's articulation of what is the *ultimate good* or happiness (*eudaimonia*) that can be achieved through practical

(political/legislative and ethical) endeavors.[2] For Aristotle, the kind of expert knowledge that one must possess and employ in order to achieve that good will be characteristic of the noble legislator-statesman who is charged with achieving it for the entire *polis*. As such, this practice of seeking the good is the noblest practical art (*technê*)[3] that one can study and exercise.

This art of achieving the good, the art of *politics*, is what Aristotle identifies as the "master" art (1094a25-b5) since it stands at the apex of all lesser *technês* and "aim[s] at what is the highest of all goods achievable by action," which Aristotle notes is generally agreed upon by "the general run of

---

[2] We need to qualify 'ultimate good' (or, *eudaimonia*) within the context of ethical/political life because Aristotle is quite adamant that the kind of good/happiness/*eudaimonia* achieved within that context is not the ultimate or *supreme eudaimonia*. The latter, as he puts it, is "the activity of philosophic wisdom" which he claims is "admittedly the pleasantest of virtuous activities" (1177a24-25). Indeed, to this "contemplative activity" "belongs […] self-sufficiency" for achieving happiness (1177a27-28). As such, "for man, therefore, the life according to reason is best and pleasantest, since reason more than anything else is man. This life therefore is also the happiest" (1178a7-9). And, again: "perfect happiness is a contemplative activity" (1178b8). Aristotle, however, goes on: "But *in a secondary degree* the life in accordance with the other kind of virtue is happy; for the activities in accordance with this befit our human estate" (1178a9-10; *emphasis* mine). What are the 'other kinds of virtues' he has in mind? He notes, for example, "just and brave acts" as well as the rest of the "moral virtues" which are "in accordance with practical wisdom" (1178b10, 15-17). What I take Aristotle to be suggesting, then, is that the kind of *eudaimonia* which is the concern of the legislator-statesman (as well as the citizens who strive to life the moral/practical life) is a sort of 'practical' *eudaimonia* – a kind of *eudaimonia* that's 'second-best' to the supreme *eudaimonia* enjoyed by the philosopher-scientist. For my purposes, however, I'm interested in his notion of the practical/ethical life (and it's 'second best' sense of *eudaimonia*) since it is the one for which he goes to great lengths to articulate a structure. Furthermore, since Zagzebski finds the Aristotelian intellectual virtue of *phronësis* to be that which is most relevant for getting us justified beliefs, and, in effect, articulates her own neo-Aristotelian aretaic epistemological account as akin to his framework regarding the practical/moral life (and *not* on his account of the contemplative life), I am more interested in pursuing that part of Aristotle's theory here in this chapter.

[3] I interpret *techne* to be 'art,' and understand it to mean a kind of expert knowledge that is employed in order to achieve some practical end. As such, *eudaimonia*, as it is the ultimate aim for the state since it is what each individual would agree to be his or her respective aim and highest good, is properly the practical aim of the legislator.

men and people of superior refinement" to be *eudaimonia* (roughly, 'human flourishing' or 'happiness')[4] even though both classes of folks, the "many" and the "wise," don't completely agree as to what *eudaimonia* or happiness actually is (1095a15-20). The disagreement here, then, is over what the specific content is that actually constitutes 'eudaimonia' (i.e., money, power, pleasure, wisdom, education, virtue, etc.).

Such disagreement regarding its content, then, makes the study of politics and ethics imperfect and perpetually incomplete. Recognizing this fact, Aristotle begins his *Nichomachean Ethics* by conceding two claims:

The first is that *the moral theory he wants to articulate here cannot and will not be exhaustive*. It is only as clear "as the subject matter admits of" since there is a limit to the amount of "precision" that the political *technê* of making one's State a good State can exercise (1094b10). So, Aristotle encourages his readers to "be content, then, in speaking of such subjects and with such premises to indicate the truth roughly and in outline" (1094b15-20). After all, the kinds of actions that "political science investigates," which are those that contribute to the success and well-being of the *polis*, "admit of much variety and fluctuation of opinion" to the extent that "they may be thought to exist

[4] "Now we call that which is in itself worthy of pursuit more final than that which is worthy of pursuit for the sake of something else, and that which is never desirable for the sake of something else more final than the things that are desirable both in themselves and for the sake of that other thing, and therefore we call final without qualification that which is always desirable in itself and never for the sake of something else. Now such a thing is happiness, above all else, is held to be; for this we choose always for itself and never for the sake of something else" (1097a30-1097b5).

only by convention, and not by nature" (1094b14-15). The study and implementation of ethics, then, does not lead to an exhaustively ultimate, absolute, or conclusive science, but an incomplete one. Yet, this doesn't mean that it is not useful (more on this later).

Secondly, this political/moral-educational undertaking is not for just anyone. The necessary prerequisite for its students is that they already have some life experience in that *polis* – i.e., knowledge of human nature such that prepares the student for the possibility of achieving full virtue. This necessarily restricts the pool of potential students. For instance, "a *young* man is not a proper hearer of lectures on political science" whether "*young* in years or *youthful* in character" (1095a1-10, *emphasis* mine). For Aristotle, the "defect" of such persons is that they are "inexperienced in the actions that occur in life" and will thus "tend to follow [their] passions" instead of the knowledge and wisdom that can only be gained through having some experience in life (1095a1-10). Such experience in preferred because it would lessen its possessor's reliance on his 'passions' and more on "desir[ing] and act[ing] in accordance with a rational principle," happily resulting in "knowledge" that "will be of great benefit" (1095a5).

Such requisite 'experience,' however, is not a protocol or formula for figuring out (i.e., reasoning) what is morally good and bad, but seems to merely provide the student with some actual *familiarity* with what has been established to be morally good and bad in the social context. Hence, the most

ready student, the one who can "listen intelligently to lectures about what is noble and just, and generally, about the subjects or political science *must have been brought up in good habits*," which suggests that his desires and actions have already been morally tempered toward what is considered good, right, and noble (1095b5, *emphasis* mine). Put another way, he must already have been enculturated in the prevailing thought of his moral community – and its superior 'wise' class, the nobility – regarding what is morally good and thus conducive to *eudaimonia*.

Of course, the relevant question to ask now is why a student of the political art must be habituated *already* in an understanding and also a practice of what is considered morally right and wrong. Why is it that the student of ethics (and, by extension of politics) must already be an experienced moral practitioner?

On one hand, there is the pragmatic reason that since he will have the task of *legislating* – performing the art of statecraft in order to make his *polis* 'good' – he must already be settled in a lifestyle of feeling and acting in the right (i.e., virtuous) ways. Because there is no universal agreement as to the content of *eudaimonia*, at least the person brought up in good habits "can easily get starting points" from which he can reason toward achieving *eudaimonia* for his community (1095b5). So, the legislator needs to be a person of good habits so that he can inculcate his *polis* in the same: "for legislators make the citizens good by forming habits in them, and this is the wish of every legislator, and

those who do not effect it miss their mark, and it is in this that a good constitution[5] differs from a bad one" (1103b1-5).

Just as with every other *technê*, then, the master art's practitioner *becomes* good at it by *doing* it, and then by extension he contributes this to his *polis* through his role as a legislator (1103a30). As such, the competent legislator is more than a mere possessor of theoretical knowledge about the good, but puts that into action by deducing judgments and creating laws that instill habits in people in order to *make* them good (1103b25). Even though studying and knowing the "collections of laws, and of constitutions also, may be serviceable" for the legislator who is charged with judging what is good and bad for his community, he who goes through "such collections without a practiced faculty will not have right judgment" (1181b5-10). *Being* good and *ordering* good requires *knowing* what is good. But as Aristotle notes, all of that requires the fundamental prerequisite of having been brought up in the good by habitually *doing* good. The role of the legislator, then, is to facilitate his *polis* in becoming good by acting good, and this is "made perfect by habit" (1103a25).[6]

---

[5] Form of government.

[6] It should be pointed out, however, that Aristotle is *not* describing the role of the legislator as dictator. This may not seem so obvious, especially when we encounter Aristotle's claim that since the majority of a citizenry cannot reason due to the pull of their individual passions, arguments alone cannot "make men good; […] they are not able to encourage the many to nobility and goodness" (1179b4-35). As a result, most people need to possess a "fear of punishment" for doing what is contrary to good (1179b14). Their actions, then, need to be "fixed by law" through the fear of sanctions that come with breaking the law (1179b35). But Aristotle is optimistic that once good actions become habitual or "customary," then people

On the other hand, being 'brought up in good and noble habits' is important because in doing this activity the student of statecraft comes to develop a love for it and can then direct his people into loving it as well. In this way the legislator is not a brutal monarch, but is one who "urge[s] men forward by the motive of the noble" (1180a5). He's not merely a legal taskmaster, a sovereign charged with governing the people for the sake of maintaining order under a continual threat of fear. Rather, he's a facilitator of excellence. Law is meant to encourage the 'many' to act in such a way that forces them to 'get a taste' of its inherent goodness, which is very important given their natural passion for the ignoble (which makes the initial practice of nobility a painful and unpleasant) (1179b30). Having mere knowledge of what is good "brings no profit" unless there is also a "desire" to "act in accordance with a rational principle" which will make practical use of that knowledge (1095a10). As such, Aristotle can strongly claim that "it makes no small difference, then, whether we form habits of one kind or another from our very youth; it makes a very great difference, or rather, all the difference" (1103b25).

So what are these excellent habits? What are these *good* actions that legislators are charged to encourage in their people? What are these 'starting points' that are prerequisite for reasoning about how to achieve *eudaimonia*?

---

who had never before "tasted" what is "noble and truly pleasant" will have indeed done so by keeping the law, and thus will experience the intrinsic beauty of good and excellent action that develops into a "kinship to excellence, loving what is noble and hating what is base" (1179b10-15; 35; 28-32).

Aristotle identifies these as *virtues*, which are dispositions that are distinctly human *excellences*. These are what "the true student of politics" has studied "above all things" (1102a5). They are excellent "states of character" that "makes the work" of its possessor to "be done well" (1106a15). Virtue "makes a man good" since it "makes him do his own work well" (1106a20). Such 'work' is what Aristotle understands to be the "function" or specific *telos* of human beings, that of flourishing in a specifically human way.[7] Human flourishing, then, is part-and-parcel of virtue. Virtue is constitutive of flourishing, that is, of being in a state of *eudaimonia*.[8] As such, virtue (practical, 'second-best') is constitutive of *eudaimonia* for individuals (as

[7] Aristotle's reason for talking about human 'function' here is simply to point out that if there is any 'excellent' way that humans particularly behave, it must be located in their capacity to reason because it is that which is distinctive of them: "the function of the human being to be a certain kind of life, and this to be an activity or actions of the soul implying reason, and the function of a good human being to be the good and noble performance of these" (1098a10-15). Alfonso Gomez-Lobo (1995) notes that 'the man who possesses reason' here refers to "someone who has the logos corresponding to a particular craft or *technê*" (p. 22).

[8] For Aristotle, however, virtue is not sufficient for achieving *eudaimonia*, and there are two specific reasons for this. First, some level of external goods is needed. Regarding the *eudaimonia* that is the concern of the statesman, "it needs the external goods as well; for it is impossible, or not easy, to do noble acts without the proper equipment. In many actions we use friends and riches and political power as instruments; and there are some things the lack of which takes the luster from happiness, as good birth, goodly children, beauty; […] as we said, then, happiness seems to need this sort of prosperity in addition" (1099a30-1099b8). Regarding the *eudaimonia* of the philosopher-scientist, he similarly notes that "the excellence of reason [...] it would seem, however, also to need external equipment, but little, or less than moral virtue does" (1178a24-25). But he also points out that even though "happiness, therefore, must be some form of contemplation," it will be such for a human being – and, so, in "being a man, one will also need external prosperity; for our nature is not self-sufficient for the purpose of contemplation, but our body also must be healthy and must have food and other attention" since a person "cannot be supremely happy without external goods" (1178b30-35). Secondly, Aristotle claims that the achievement of *eudaimonia* will be the product of virtue plus the necessary amount of external goods *over one's life span*: "For there is required, as we said, not only complete virtue but also a complete life, since many changes occur in life, and all manner of chances, and the most prosperous may fall into great misfortunes in old age, as is told of Priam in the Trojan Cycle; and one who has experienced such chances and has ended wretchedly no one calls happy" (1100a4-10).

citizens/subjects) and hence for societies. For us, then, possessing and exercising these excellences contribute to making us *become* good, not only in terms of our actions, but also in terms of our passions as well (1106b25).

One of the most important and interesting aspects of Aristotle's notion of virtue is that even though he describes it as one side of a binary – its opposite being a moral *defect* (or what has been commonly referred to as a 'vice') – the defect itself can take form as either of two opposite extremes to that virtue. Virtue itself, then, is actually the intermediate (or mean) state located between two defective states of deficiency and excess.[9] As such, "virtue must have the quality of aiming at the intermediate" (1106b15) where 'intermediate' is to be understood as "not in the object but relative to us" (1106b5-10).

To illustrate, consider that feelings such as "fear and confidence and [the] appetitive and anger and pity and in general pleasure and pain may be felt both too much and too little," and so what is "characteristic of virtue" is to "feel them at the right times, with reference to the right objects, towards the right people, with the right motive, and in the right way" which is "intermediate and best" (1106b15-20). And so, feeling them in too much of a

[9] Of course, it's not the case that all virtues, for Aristotle, are at a mean between extremes and defects. Aristotle notes that justice, for instance, is a virtue that is not located between some excess or deficiency (see Book V of the *Nichomachean Ethics*).

way (excess) or in too little (deficient) of a way would be contrary to virtue (1109a20).[10]

Of course, this also applies to actions. For instance, the expression of the virtue of courage, which is the manifestation of bravery in a fearful situation, is "destroyed by excess and defect" (1104a25). A deficient expression would not be courage, but cowardice. Expressed in excess, it would be something like haughtiness or arrogance. The character of a virtue, then, is that *it always is an expression of the proper (intermediate) state* relative to the particulars of the situation that calls for it – whether a passion such as anger and joy, or an action such as generosity, moderation, and courage – or else it isn't a virtue.[11]

This implies a crucial distinction, I think, that may help to understand what Aristotle may be trying to get at here (even though Aristotle doesn't make this distinction himself). It is one thing to talk about *exercising* a virtue – that is, exercising a character trait that one possesses – but it is something different (yet related, obviously) to talk about how that virtue ought to be particularly *manifested* in a particular situation. On Aristotle's account, since

---

[10] Passions and actions are *correct* and are thus the expression of virtues when virtues strike a mean with respect to them. As Aristotle claims, "moral virtue is a mean […] between two vices, the one involving excess, the other deficiency, and that it is such because its character is to aim at what is intermediate in passions and in actions" (1109a20).

[11] Of course, I'm simplifying this a bit. For Aristotle, while courage is a state that strikes a mean between fear and confidence, its opposite vice can take form in either of four different ways: one vice of excess is too much fear but another vice of excess is too much confidence; another vice of defect is too little fear and a fourth is too little confidence (See *Nichomachen Ethics*, Chapters 6-9 of Book III).

the virtuous person *exercises* virtue reliably and consistently, this person will know *when* to do this. But knowing when to do it seems conceptually different and distinct from knowing *how* to do it (even though it would seem that even if Aristotle accepted this distinction I'm making here, he would probably argue that the truly virtuous person will be reliable at both). This is merely to say that virtue is manifested in different ways respective to the nature and particular details of a situation that calls for its exercise. One implication of this, then, is that given the complexity and distinctness of the many different circumstances of life, virtues are rarely ever manifested in the same way for every circumstance. Nor is there is a sort of explicit 'one size fits all' character to virtue, even though Aristotle would claim that any virtue will be the same virtue for one person as for another (assuming, of course, they are virtuous) even though it would seem that the way those virtues will be manifested in different situations will be different respective to the particulars of those situations.[12]

To illustrate this, consider how my virtue of generosity might be expressed in two different sets of circumstances – toward my child and toward a local charitable organization that collects food for the homeless. How should

[12] As Richard Kraut (2014) notes, for Aristotle "the intermediate point that is chosen by an expert in any of the crafts will vary from one situation to another. There is no universal rule, for example, about how much food an athlete should eat, and it would be absurd to infer from the fact that 10 lbs. is too much and 2 lbs. too little for me that I should eat 6 lbs. Finding the mean in any given situation is not a mechanical or thoughtless procedure, but requires a full and detailed acquaintance with the circumstances."

I appropriately manifest generosity in each of these situations? As I think about this, I'm reminded of an old adage – 'give a man a fish he will eat for a day; teach him to fish and he will eat forever' – that gets its intuitive thrust by reminding us that it's important for an individual to have learned the skills for self-preservation and self-sustenance because it is more sustainable than being utterly dependent on others. This seems like a pretty universally accepted fact of human life and may even apply to every sphere of life, not just that of physical nourishment. After all, a healthy society is usually considered to be one where, at the very least, people have the freedom and skills to provide for themselves, think for themselves, feed themselves, clothe themselves, etc.

Yet, even though this may indeed be true, does this mean that my generous response to the needs-requests from the local charity *should* take the form of my showing up to 'teach' everyone how to 'fish,' that is, how to secure their own sustenance rather than relying on others for it directly, even if I consider this to actually be a *generous* offering of the life-skills that I've learned 'the hard way' and now want to share with these others? Probably not. Such behavior would almost certainly be considered rude, insensitive, and I would probably be appraised as a condescending jerk. No doubt, there may be a more proper way to go about helping those who need to transition into a more self-sustaining livelihood. But that role is appropriately one for professionals to fill – those who have the essential tools and training to help these particular kinds of dependants realize that goal. It can't be achieved by my twenty-minute

oration on 'how ya'll need to go out and get a job.' My appropriate response to the call of exercising my virtue – that is, the correct manifestation of my virtue of generosity – then, will be in offering what I can appropriately, relevantly, and *suitably* contribute. And in this particular case, being generous will be cashed out in terms of giving some of my own tangible resources – money, canned goods, my old car, etc. Exercising generosity, then, will take the form of this particular *manifestation* since it is the proper one for this unique situation.

However, there may be other situations where generosity is actually best manifested by waxing eloquently an oration on 'teaching how to fish' rather than just giving them away.

Consider, for instance, the particular responsibility I might have as a parent of teenage children. It would seem that *in light of* my particular *role* as a father to these children (which would seemingly be the relevant factor that defines this particular context now), being *generous* toward these people (i.e., my children) might involve other things rather than just giving them stuff. Indeed, it many include non-tangible things such as imparting to them some of the wisdom and life-skills I have gained from my own experience about the kinds of things that contribute to successfully fulfilling one's life and vocational goals that by extension allows one to effectively provide for their own sustenance. Maybe the most generous thing I could ever do for them would be to withhold 'freebies' and replace them with the motivation and wisdom that

encourages them to obtain those things through their own efforts. As such, articulating to them in the course of a few years of family suppertimes what has taken me the course of my lifetime so far to figure out would be very generous indeed. A quick perusal through the 'self-help' and 'leadership' sections at the local bookstore indicates that such wisdom is worth around $34.99 per copy. Passing this on free of charge to my children, then, might be one of the most generous things I could ever do for them, specifically in terms of securing their own sustenance, but especially in terms of what it will do for their own self-esteem.

In short, then, this particular manifestation of generosity (i.e., withholding tangible stuff in order to teach how to get it for oneself) might be very appropriate in the particular situation regarding my children, *while at the very same time* it would be very inappropriate in the situation where the recipients are those who need the services of the local homeless charity which has called upon me to give tangible stuff.

What I think this helps to illustrate is Aristotle's point that the *way* a particular virtue should be manifested by a virtuous person in a particular situation will not be universal across all situations. As Aristotle himself puts it, the virtuous person "seeks the intermediate and chooses this – the intermediate not in the object, but relative to us" (1106b6-8) whereas " the intermediate relatively to us [is] that which is neither too much nor too little – and is not one, nor the same for all" (1106a28-32). Another way of putting this may be

that the appropriate *manifestation* of virtue (what the virtue should look like when it's exercised) depends on the particulars of the *context* in which that virtue will be exercised (I'll elaborate on this more in the next section).

So here's the million-dollar question: how do I accomplish this? How do I figure this out? I may know that I need to exercise generosity in a particular situation, but how is it that I come to know *how* to properly and appropriately do it in a way that is suitable for that particular situation? What's a generous 'tip' in a truck stop cafe as opposed to a high-end French restaurant? What are the guidelines/rules/standards that I need to appeal to in order to do this correctly?

One of the interesting (and for some, disappointing) aspects of Aristotle's ethics is that there are no such rules, no such guidelines.[13] He offers no strategy or procedure or formula for determining the intermediate. He simply seems to defer to some notion of an *exemplar*:[14] the virtuous person is so constituted that he is able to determine *how* a virtue should be manifested *in* a particular situation and do so in a way that it appropriate or 'relative' to that

[13] Richard Kraut (2014): "So far from offering a decision procedure, Aristotle insists that this is something that no ethical theory can do. His theory elucidates the nature of virtue, but what must be done on any particular occasion by a virtuous agent depends on the circumstances, and these vary so much from one occasion to another that there is no possibility of stating a series of rules, however complicated, that collectively solve every practical problem."

[14] Linda Zagzebski (2010) argues that "basic moral concepts are anchored in exemplars of moral goodness, direct reference to which are foundational in the theory. Good persons are persons *like that*, just as gold is stuff *like that*" (p. 51). As such, moral concepts are "defined via indexical reference to a paradigmatically good person" (p. 54). I will discuss this at more length below.

respective situation *in the way that the* "man of practical reason" *or virtuous person would do it* (1107a1).

Of course, this kind of 'figuring out' is not without difficulty. Aristotle concedes as much: it "is not easy to determine by reason" what is the intermediate since "the decision rests with perception" on the part of the agent, and given how he perceives the situation, "he will incline sometimes toward the excess, sometimes towards the deficiency" (1109b23-27). So, generosity may sometimes be manifested best near the vicious fringes of excessive lavishness (such as when one bequests one's multi-million dollar estate to a charitable organization that is able to handle such a large gift and make good use of it without difficulty), and sometimes it's best manifested close to the edge of deficiency (such as in a $25 rather than $500 birthday gift for an 8-year-old). The virtuous person is called upon to employ his own judgment, and of course, it is assumed that he will do this in a way that gets it right since, by definition for Aristotle, a 'virtuous' person necessarily must!

The important point to note here is that even though Aristotle articulates no exhaustive set of rules or standards or procedures for how to do this kind of moral decision-making work, the student should not come away from this discussion with the belief that discovering the intermediate is a completely relativistic and exception-rife enterprise. As Richard Kraut (2014) notes, "Although there is no possibility of writing a book of rules, however long, that will serve as a complete guide to wise decision-making, it would be a

mistake to attribute to Aristotle the opposite position, namely that every purported rule admits of exceptions, so that even a small rule-book that applies to a limited number of situations is an impossibility."

For Aristotle, virtue, which is located in "the mean relative to us," is "determined by a rational principle and by that principle by which the man of practical wisdom would determine it" (1107a1). The virtuous person possesses *practical wisdom,* which enables him to employ his *rational capacities* in such a way that allows him to correctly *determine* what is the appropriate way to manifest virtue in a particular situation.

Let's break these concepts down.

Practical wisdom (*phronësis*) is the intellectual virtue that equips a person with the ability to "deliberate well about what is good and expedient for himself" which is to say that he "calculates well with a view to some good end" (1140a20-30). Aristotle is very specific about this: the person possessing *phronësis* doesn't deliberate about good ends since he already knows what the "good life" is, so the objects of his deliberation will be those things that contribute to reaching those good ends (1140a29-30). This person has the correct perspective about the good. He just has to *deliberate* about the *means* to achieve that good:[15] "we deliberate not about ends but about means"

[15] A note of clarification is appropriate here. In articulating Aristotle's 'doctrine of the mean' above, I have purposely avoided referring to virtue as a 'mean' but rather as an 'intermediate' for the sake of avoiding ambiguity as well as the possibility of unintended equivocation with 'mean' understood here as the *way or method towards achieving some goal*. As such, my use of

(1112b10). For example, medical doctors don't deliberate about whether or not they "shall heal," nor do statesmen deliberate about whether or not they "shall produce law and order," but what they all deliberate about is by "what means [their respective goals are] to be attained" (1112b10-15). As such, *phronësis* is "concerned with things human and things about which it is possible to deliberate" since "no one deliberates about things invariable, nor about things which have not an end" (1141b5-10).

The person with *phronësis,* then, judges according to a 'rational rule' (*orthos logos*) which Aristotle identifies as a reasoned or right/correct account. He explicitly claims that striking the intermediate is "determined by the dictates of *right rule*" (1138b15) which can only be accomplished by the person with *phronësis,* with the "*reasoned* and true state of capacity to act with regard to human goods" (1140b20). *Phronësis*, then, is the excellence of rationally figuring out the means to achieving practical ends and its performances are well-accomplished deliberations that result in correctly reasoned accounts.

Of course, the mark of a reasoned account is that it is *justifiable*, that is, that its accounting is *articulable*. Aristotle points out that it's not enough to say that virtue is an excellence that accords with a right rule, but rather, "we must go a little further" and understand it as "the state that implies the *presence* of

---

the term 'mean' or 'means' will specifically be intended to represent this latter idea throughout this chapter.

the right rule" (1144b25). As such, that virtue carries with it an inherent *orthos logos* implies that it is something that its possessor can justify.

The question that now arises is this: justify against what? After all, in order to justify something, there needs to be some model or procedure that the justifier points to as the standard against which her justification is determined, but we've already said that no specific/universal procedure really exists for Aristotle.

Although I will discuss this more in the next section where I discuss at length what I think is a plausible contextualist interpretation of Aristotle's basic ethical framework here, allow me to speculate here that it would seem that the *orthos logos*'s justification might plausibly contain a *comparison* to the relevant exemplar (i.e., the virtuous person) as well as some *contrast* to other situations (i.e., other circumstances which call for the exercise of the same virtue, but not in the same way or manifestation as the one in question). As such, the virtuous person grasps why it is that the particular intermediate she strikes is correct. After all, she would seemingly be *aware* of all those features (her own as well as those of the situation she is in) since her "decision rests with perception" (1109b23).

So, what we get here is an understanding of *phronësis* as that which equips the virtuous person to determine the appropriate intermediate and, hence, models a sort of decision-making procedure that is to serve as a sort of

standard to be imitated, as cohering with that of the moral exemplar. Richard Kraut notes that

> The intermediate point that the good person tries to find is "determined by logos ('reason,' 'account') and in the way that the person of practical reason would determine it" (1107a1–2). To say that such a person "sees" what to do is simply a way of registering the point that the good person's reasoning does succeed in discovering what is best in each situation. He is "as it were a standard and measure" in the sense that his views should be regarded as authoritative by other members of the community. A standard or measure is something that settles disputes; and because good people are so skilled at discovering the mean in difficult cases, their advice must be sought and heeded. (2014)

In sum, then, the virtuous person possesses *phronësis*, the exercise of which empowers her to deliberate about *how* a particular virtue should be manifested across the spectrum of the various circumstances in life (and hence the various ways that virtue could be manifested) that call for it. We must be careful here because Aristotle is clear that the virtuous person never chooses whether or not to exercise a virtue. The virtuous person has no such choice since his virtue flows from a "firm and unchanging character" (1105a30), which suggests that he is a perfectly reliable exerciser of virtue.

Along the same lines, Aristotle claims that because the virtuous person is "truly good and wise" he "always makes the best of circumstances" (1101a1-5), which suggests that he will know when and how to appropriately manifest virtue in every circumstance.[16] So, since *eudaimonia* is the highest good and

[16] These claims are the particular targets of the 'situationist critique' that will be discussed in the next chapter. In short, this critique claims that such examples of virtue are non-existent since it can be shown that some examples of virtuous action result not from the person's character, but in response to features of the situation itself. One way that I will try to respond

ultimate end (although we've recognize that, for Aristotle, he seems to talk about two kinds of *eudaimonia* and we're concerned more with that kind which is the goal of ethical/practical virtue), and since virtue is a necessary condition for reaching *eudaimonia*, the person who is virtuous (and also has enough external goods) is, by definition, *eudaimon*.

As such, the particular kind of moral reasoning or deliberation that contributes to one's achievement of *eudaimonia* is that which allows the agent to determine *how* virtue is to be manifested, i.e., what is the "intermediate" that is "in accordance with the right rule" for a particular situation (1138b30).[17]

Consider one way that Aristotle illustrates this. He claims that the only proper objects of deliberation are those things "less exactly worked out" such as the master *technê* of politics and also the *technês* of medicine and navigation (1112b5). Why the latter two? Presumably, in Aristotle's time, these were less exact than they are today given our vast array of medical imaging and diagnosis-testing technologies at our disposal, as well as satellite assisted global positioning systems, digital compasses, radar, and sonar. Aristotle and his contemporaries didn't have anything remotely close to this kind of technology.

---

to this criticism is by noting that on an Aristotelian scheme features of the situation are indeed important parts of virtuous action.

[17] While this understanding or interpretation of the role of choice in Aristotle's ethics may be controversial, it seems to me to be the best way to understand, on the one hand, Aristotle's claims that for the virtuous person "actions and states of character are not voluntary" (1114b30), with his further claim that a virtuous person must "have knowledge, secondly he must choose the acts, and choose them for their own sakes" (1105a30), as well as his claim that "virtues are modes of choice or involve choice" (1106a1) and that virtue "is a state of character concerned with choice, lying in a mean, thus being determined by a rational principle" (1107a1).

So, consider Aristotle's claim that I noted earlier how that the master *technê* of politics – the most noble one given its crucial role in governing one's practical life as well as one's *polis* – cannot reach any kind of precision or exactness regarding how it works, and so, we must "be content" with the truth that is only "roughly and in outline" (1094b15-10). Of course, its constitutive state – *eudaimonia* – is securely positioned and permanently fixed. But a clearly worked out procedural method of reaching this state isn't. This is where deliberation is called for.

Medicine, for Aristotle, is a deliberative art akin to the master art (1112b1) because of what it requires from its practitioner: he has to figure out how to best achieve its goal (i.e., health) since the ways of attaining health are varied and crucially depend on aspects of a particular case in question. Presumably, the physician must take a lot of things into consideration when treating a patient's illness such as her age, weight, past illnesses, current medications, etc., before the best treatment for her can be prescribed. So, even though health is the indisputable fixed goal of the medical art, the means to achieving health aren't so much. They are not universally applicable (in terms of their every detail) for every patient in every situation.

Similarly, then, to possess and exercise virtue (and, thus, be *eudaimon*) is the indisputable goal of the master art. But the specific way to always express and manifest virtue is non-universal across every situation and circumstance. As Aristotle notes,

> Now, that we must act according to the right rule is a common principle and must be assumed […] but this must be agreed upon beforehand, that the whole account of matters of conduct must be given in outline and not precisely, as we said at the very beginning […] matters concerned with conduct and questions of what is good for us have no fixity, any more than matters of health. The general account being of this nature, the account of particular cases is yet more lacking in exactness; for they do not fall under any art or precept but the agents themselves must in each case consider what is appropriate to the occasion, as happens also in the art of medicine or of navigation. (1103b30-1104a10)

The conclusion Aristotle reaches, then, is that this particular art – the art of achieving the good life for one's self and one's society – is an inexact and incomplete one.

Of course, in one way this makes perfect sense. If there were some precise formula for doing this sort of work it would no longer be a *technê* but an *epistémë*, no longer an art form, but an exact science. But it doesn't reach this kind of high epistemic status for Aristotle. This *does not* mean, however, that the moral judgments reached (i.e., the intermediates settled on and manifested) in particular cases are not certain. It's just that the means of reaching those judgments aren't as methodologically precise or defined (and hence not as epistemically definite as a 'hard science,' for instance).

But this poses a crucial question: does Aristotle's ethics necessarily rely too much on relativity for it to be useful? Is it not objective enough to be plausibly normative?

Consider Kantian ethics, for instance. Although highly contentious, Kant at least gives a sort of rubric or formula (i.e., the categorical imperative)

by which a maxim can be tested and hence whether the action that would follow that maxim should be prescribed or prohibited. Or, consider a Utilitarian-styled ethics that gives us the principle that a correct moral action is one that maximizes the most amount pleasure or happiness for the greatest amount of people. Aristotle gives no such rules or principles to us. We have no objectively *principled* way to approach ethics on his account. All we have is a notion of a virtuous person who possesses practical wisdom, a *phronimos* who has been reared in good habits and has fallen in love with virtue so much so that his state of character has achieved such a high moral plane. He reliably responds to situations he encounters in the right way given the insight he has in those situations, and thus can form a reasoned account of why intermediate he struck was indeed the correct one.

As such, what Aristotle *doesn't* give us are *principles* of moral *reasoning*. However, what he *does* give us is the concept of a *person* who is a good moral *reasoner*, and out of this we can define what are good moral *actions*: they are the actions that are done in the way that a virtuous person would do them. As Aristotle explicitly notes, "actions, then, are called just and temperate when they are such as the just or temperate man would do" (1105b5). All that Aristotle really gives us, then, is this description of a *model* or moral exemplar accompanied with the assumption that in order to live well, that is, to live according to virtue, you must imitate this model (in terms of the ingredients or characteristics of this model such as virtue, *phronësis*, good habits, the right

external goods, etc.). Such a model, then, is the *exemplar* which is to be imitated in order to correctly make one's moral decisions.

3.3. A Contextual Reinterpretation of Aristotle's Ethical Framework

In this section, I would like to consider some of the contextual overtones that I think are implicit in the discussion of Aristotle's moral theory in the previous section. As such, I want to provide something of a neo-Aristotelian account of an aretaic ethical framework based on the basic rubric outlined in the above account. As such, there is a way, as illustrated by the somewhat radical re-interpretative work of Peter Simpson and Martin Tweedale, to consider that account as one that not only utilizes *macro* and *micro* contexts (although not in so many words, obviously), but also manifests an important interplay between them in a way that has affinities to the kind of contextualism that I discussed at §1.2.

There, I argued that the notion of 'attributor factors' – those standards that epistemic agents must satisfy in order to make true knowledge attributions – are collectively determined by the epistemic community at large. As such, this social context – the large collective that sets epistemic standards and norms – is what I have termed the *macro*-context. But also, there are 'subject factors' – the aspects of an individual agent's situation – that are the facts about the situation that tell us (1) what the rules/norms are for that situation which have been set by the broader social (macro-) context and also (2) whether or

not the individual has actually met those norms. If she has, then she can correctly be attributed with having some sort of positive epistemic status such as what the community wants, needs, and requires of us. The situations in which subject factors are exercised and discovered are what I have termed *micro*-contexts. Each micro-context is unique just as every situation in life is unique. And so, the relevant degree to which any positive epistemic status can be attributed in each micro-context will understandably be non-uniform among all micro-contexts.

So how might we think about the rubric Aristotle gives us as plausibly being re-interpreted as implying a sort of contextualism that I have in mind?

First, consider that there is some sense in which norms are collectively determined as such by the macro-context: this is embodied in the ultimate moral model, the exemplar of virtue. This is the standard by which the relevant object (here, *morality*) is to be achieved.

Second, there are micro-contexts, those various situations and circumstances in which good is achieved by virtue being exercised in them. Plausibly, this may sometimes require the exercise of the very same virtue in many different situations although the way in which the virtue is manifested (i.e., the intermediate that is struck) will look different in each of those situations. So, even though the norm or standard – the exemplary virtuous person – has been set to look a certain way at the *macro* level, it is at the *micro* level that we look for the 'subject factors' in order to see whether or not the

subject has indeed lived up to that standard, that is, if she has performed in a way consistent with the socially-sanctioned macro-model.

First, then, consider why we might plausibly interpret something like a ***macro*-context** as implicitly taking place in Aristotle's basic ethical framework. His entire scheme relies upon a very specific and narrow understanding of the exemplary virtuous person based on the macro-contextually sanctioned understanding of what this kind of person looks like. As such, we don't really find an argument about what the exemplar ought to look like, but what we actually get is a description of the phenomena that Aristotle and his readers would already be familiar with in terms of the "paradigm of human excellence" which for them, as Alasdair MacIntyre claims, is the "Athenian gentleman"[18] (1981, p. 182).

One reason, then, why Aristotle may be conceding that his moral theory is incomplete and hence articulable only in rough outline is because there exists for him and his contemporaries (both in time and locale) a very non-controversial standpoint that the *noble* are *virtuous* and that this latter appraisal can only truly be defined in terms of the former.

This also may help to answer some of the questions we are left with, such as why a virtuous person must have first been brought up in good habits?

[18] Of course, the sort of exemplary class that Aristotle has in mind wouldn't have been described by him as 'gentlemen' since this term is an 18th century British one. But, it may serve as a decent 'catch-all' in English for the various classes of citizen (not all of them) that would have been held in esteem as moral and political exemplars because these are the ones who have 'been brought up well.' And this may precisely be the reason why MacIntyre uses it. More on this later.

Why must he must be forced by law to *act* good so that he will come to love what is good (i.e., virtue) and then actually *become* good? Why is the ultimate principle for acting in accordance with morality not a principle at all, but a *person* whom one is supposed to mirror? Furthermore, why this particular exemplar and not another?

In his radical critique of Aristotle's ethics, Peter Simpson (1997) argues that the answer to these questions as well as "whence [Aristotle] derives the virtues?" and "what other justification does he give in their defense?" is a "very simple answer: from common opinion" (p. 248). For Simpson, "the virtues Aristotle lists, and the descriptions he gives of them and of their possessors, are taken from the common experience and opinions of citizens of the day" (p. 248). But of course, this would not be the common opinion of the 'many' since, as Aristotle himself noted, they would consider things like sensual pleasure to be the contributing factor to human flourishing, not virtue (1095a20).

So, if this is the opinion that Aristotle is relying on to corroborate or justify his particular ethical scheme, who are the 'wise'? They are those "called by Aristotle, 'the beautiful and the good'" which Simpson argues our English words "gentlemen," "nobles," and "aristocrats" signify adequately well (1997, p. 248). Although such terms in English might not perfectly represent the different ruling classes of Athens at the time of Aristotle, the point is that those who would've been the modeling examples of *phronësis* – possessing

intelligence and prudence – would've been men who also held these statuses. And one way to maybe think about the reason why they would have held such statuses is because of the privileged life they lived in not having to worry about providing an income, and thus having the time and energy to devote to other endeavors, especially that of guarding and supporting the State (and hence being blessed with those 'external goods' or 'equipment' that Aristotle also noted as a necessary condition for the achievement of *eudaimonia*).

An Athenian 'noble' (*eugenes*) would've been a very wealthy landowner, someone whose mass of wealth and property would have allowed him to live very well on that bounty alone (without significantly diminishing it) and would not have had to concern himself with bringing in any sort of revenue by any other means. An Athenian 'aristocrat' (*aristokratía*) would've been a landowner whose estate would have provided employment for others who worked his land for him and hence the profits of its fruits would've been the income that landowner would've received. As such, whether an aristocrat or a gentleman, the point is that someone born into one of these positions would indeed have been brought up in such a way as to have learned the kind of 'good habits' that Aristotle noted was the necessary prerequisite for becoming a good statesman since they would have had the time and means to do so.

At the same time, however, there would have been a sort of duty to do this, that is, to serve – to assume one's responsibility of managing the State as a proper response to having been part of the privileged classes in Athens. This

would've especially suited the privileged person since they would have already possessed some experience at managing themselves as well as their estates.[19] As such, men of these statuses would've made up the ruling class. (And, of course, we can assume that the kind of Athenian citizen that would've even had the leisure time to spend with Aristotle and read his works – i.e., the target audience Aristotle is writing for – would've been just this kind of person).

So, even the kind of wisdom (*phronësis*) that is necessary on Aristotle's particular brand of moral reasoning for determining the appropriate intermediate would not really seem to be something that everyone could do since "not everyone's intuitions count as instances of prudence" but only those "of the virtuous" (Simpson, p. 249). And so, we end up with this circle: the virtuous are the noble/wise, the noble/wise are the virtuous. And this circle is not a problem for those who are in it, but it *is* "a problem for everyone else" since they need the forceful sanction of the ruling class in the form of laws that will make them good in order for them to actually become good (p. 249).

On this account then, virtue is unintelligible for the person who has not been enculturated by the method of the macro-contextual exemplar. That 'being brought up in good habits' is a necessary condition for being virtuous is just another way to say that one cannot be virtuous – cannot live up to the *macro*-determined moral standard – unless one has been macro-contextualized

[19] I am indebted to Bruce Hunter for helping me parse out the various classes of Athenian citizens and why this would be relevant to Aristotle's understanding of the *phronimos*.

in one's moral community. Aristotle claimed that such good habits are necessary for getting the 'starting points' that one needs for successful moral reasoning. But one might, then, argue that these 'starting points' are really just basic moral concepts that are macro-contextually determined. Once a person has these starting points – that is, once one has been macro-contextually enculturated – then she can extend her moral-reasoning ability and capacity in the various situations (micro-context) she encounters.

What we're now seeing, then, is why Aristotle's ethics is incomplete, that is, why it is devoid of the kind of content that gives us procedural criteria or rules for making ethical decisions, and yet it still is able to account for why a virtuous person gets it right when she does. The truth about particular virtues in terms of what the 'intermediate' is supposed to be in a particular situation is not "shown by theory" but is "perceived by prudence" and then praiseworthy when it is the kind of thing the exemplar would do (Simpson, p. 250).

Furthermore, consider those things that Aristotle tells us are blatantly wrong without question or qualification: the passions of "spite, shamelessness, envy, and in the case of actions adultery, theft, murder" (1107a10). These have no intermediate. They are always immoral. But why? Aristotle offers no principled way that determines their 'badness,' but rather, seems to operate off of assumption that his hearers will already agree with him about these. Presumably, this assumption may simply be because such precepts are automatic and indisputable norms given the way that human beings are

structured, which would be discovered and then prescribed by the 'wise' class which, again, seems to be the influential arbiter of what is morally 'excellent'. These passions and actions which are noted as unqualifiedly bad, as well as virtues such as "justice" and "courage," are all mere "givens" for Aristotle (p. 250). They are not argued for. They are just conclusions. Simpson's radical conclusion from this is that "Aristotle's ethics is an ethics of and for the gentlemen. It is prejudice, not philosophy" (p. 250).

Now, whether or not one agrees with Simpson's conclusion (and moreover his analysis altogether), however, shouldn't distract us from the insight that he does offer as to why it may be that Aristotle places his moral theory within a political one. The moral scheme Aristotle articulates is not something that people will come to discover naturally, but only through arduous discipline. And, so, there arises the question of how to satisfy one's social responsibility to do what it takes to provide this discipline. Furthermore, *cultivating a love* for what is noble and good (in the way Aristotle describes it as a necessary condition for being virtuous) is not something that the 'many' will be able to achieve on their own. Hence, Aristotle argues for the proper (and noble) role of legislative force.

But realize what this means. For Simpson, Aristotle's notion of moral wisdom cannot exist "without architectonic legislative prudence" (p. 255). Aristotle's picture, then, is one of "moral and political hierarchy" where the rulers who "possess virtue in the highest sense and prudence in the highest

sense" are located at the top, and "below them in descending orders and dependencies come the various kinds of good citizen and good members of households" (p. 255). What Aristotle gives us, then, is an ethics that "is the product of a good regime" (Simpson 1997, p. 256).

Consider how, I would suggest, that Martin Tweedale also weighs in on what I would call this macro-contextual aspect in Aristotle's ethics, although Tweedale is more charitable to Aristotle's interpretation of his own ethical framework than Simpson may be. As Aristotle argued, a person can develop the master art only if he has attained the proper 'starting points' which arise from one having been habituated into good habits (1095b5). As Tweedale notes, this

> requires that the discussants come with a character that is already pretty darn good, otherwise no agreement will be found on starting points, which I take to be judgments about whether certain choices that people have made, whether these are drawn from history or legend, are noble or base. In other words, the discussants have got to agree on and recognize as noble or base the more blatant cases of such behavior (2007, p. 3).

What this may suggest, then, is that Aristotle assumes that the principles of his ethical inquiry are things that people agree to – or, as I should point out, that the *people who matter* (i.e., the wise class) agree to. And this seems to be Tweedale's point. For Aristotle and his followers, there would be no use in arguing with types of people such as "Conrad Black" or "Ayn Rand" or maybe even Bernie Madoff, for instance (p. 4). Why? Because while some communities might admire these sorts of (im)moral exemplars, these are not

the kinds of people whom Aristotle's particular social context would have held in high esteem (and, hence, wouldn't have been proper models of a moral exemplar for him/them). But in holding up as a model the particular macro-contextually influenced exemplar that he does, Aristotle "enthrone[s] the intuitive judgments of the dominant culture or the majority and dismiss all those who have fundamentally different intuitions" (Tweedale, p. 4). So, what seems to be the case is that since there is some sense of agreement at the macro (ruling class) level as to what the correct exemplar is, which may be to say that this particular community authoritatively defines what the good life is and what being good truly means, it may be thought that this is enough to make it normative for the entire State. Hence, this becomes the model of moral practice that is instructive and authoritative for everyone in that State. Of course, Aristotle doesn't argue this last point explicitly. And I don't want to attribute it to him nor to his own interpretation of the basic framework of morality that he outlines for us. But maybe this is just implicitly assumed given the normative role played by the exemplar that is influenced (constructed?) by the wise class. (More on this in §3.5 where I discuss a possible foil to this particular interpretation of Aristotle).

So, we see how the *macro* and *micro* may be *related* here whereas the *macro* determines the *micro*. The broader social context plays a role in determining what is the right way to live in general (and what is the moral thing to do, in particular) through the standard-bearer (exemplar) which is the

rule by which persons are to act in particular micro-contexts even though that *exemplar* is such because it exhibits all of those traits deemed beneficial and moral, as Simpson noted earlier, by the common opinion of the day. The standard for moral reasoning, then, is the exemplar.[20]

As such, the exemplar would be somewhat of a social construction. It is identified by the larger social context in light of what that society values and needs. But if we're uneasy attributing this to Aristotle (and, to be charitable to him, we have to emphasize the point that this sort of interpretation of the historical Aristotle is *not* his), maybe we can make the same point by calling it a 'social recognition' of values and standards at a broad level that are determined in virtue of the needs and interests of the broader society. And, then, it's up to the individuals within that larger social context to apply them more specifically to their own circumstances that they find themselves in.

To conclude this section, then, I have offered Simpson's and Tweedale's somewhat radical re-interpretation of Aristotle's ethical theory in

---

[20] One interesting question that this raises is whether or not it's plausible to say that a person can have an *orthos logos* without a methodological procedure for moral reasoning but, instead, with an exemplar or model serving as the standard and hence as the justification for the account that will be made. One response might be that Aristotle's notion of an exemplar qualifies as a sort of moral reasoning procedure in light of the two questions the agent must answer when trying to match her behavior in accordance with that of the exemplar: (1) How would action x compare to what the exemplar would do in a particular situation s^2? (2) How is a particular manifestation m^2 of x the appropriate intermediate particular to s^2 as opposed to how a previous particular manifestation m^1 of x was the appropriate intermediate in situation s^1? One might argue that the asking and answering of these questions is the procedure itself. Put another way, this procedure is as I described it earlier in this chapter: the agent must *compare* her behavior with the exemplar and then *contrast* its manifestation with other similar situations in order to figure out the appropriate way of responding to the current situation at issue.

that they serve my purposes by highlighting the significance of a *macro*-context and *micro*-context (as I have understood it in the previous chapters albeit here regarding ethics rather than epistemology) as possibly being something that is motivating the claims Aristotle makes. As such, the important part of this sort of interpretation would be how both kinds of contexts interact – the macro is that which sets the (moral) norms for various micro contexts where those norms are embodied in the notion of an exemplar. Of course, Aristotle's scheme may not be the only virtue ethical account that arguably employs this sort of paradigm, and this may be unwelcome news to those who share Simpson's radical distrust of any ethics where elites 'call the shots.'[21]

Another such account may be that of the neo-Aristotelian Alasdair MacIntyre. While the above arguments are controversial among many Aristotle scholars, we nevertheless get a similar kind of micro/macro distinction from the work of MacIntyre. He offers a more promising re-working of Aristotle's ethical framework (from what has been proposed in this section), yet he offers somewhat similar way of thinking about this in terms of what it means to *practice* virtue where such practice implies that we explicitly respect the macro-context from which virtue obtains its legitimacy. His ideas are the subject of the next section.

[21] Such as how Peter Simpson (1997) views Aristotle's ethics: "Virtue, both in its subordinate and its highest forms, is the product of a good regime, and a good regime is one where gentlemen are dominant" (p. 256).

3.4. Virtue's Inescapable Plight from the Contexts

In discussing virtue (whether moral or epistemological), one of the more proximate problems one faces is coming up with reasons for choosing which aspect of virtue one wishes to discuss (or emphasize), and hence which 'theory of virtue' one will want to appeal to. The field of virtue theory is legion. In its entirety as a corpus, it is less like the uninterrupted strands of uniformly cut wire that make up the cables holding the spans of the Ambassador Bridge connecting Detroit to Windsor, and more like the irregular and short strands of hemp that lock together to form the frame of the footbridges of old: it is strong and useful even though upon closer inspection we see that the structure itself is complex, varied, and diverse.

In this section, I want to consider a suggestion of the neo-Aristotelian Alasdair MacIntyre who retains a basic strand of Aristotle's moral framework (in that he follows in that tradition, however, without buying the whole thing) which he develops more fully and interestingly. The point that MacIntyre makes is that one aspect of virtue that must not overlook is that fact that they are *practiced*.

MacIntyre (1981) notes that even though the virtue-theory tradition is marked by different understandings of virtue, there nonetheless is a "core concept" within that tradition that gives it its "conceptual unity" (p. 186). Each theory may exhibit a different understanding of 'who' the paradigmatic virtuous person (*phronimos*) or exemplar is, whether it's Aristotle's "Athenian

gentleman" or Homer's "warrior," for instance (p. 182). But in positing such a paradigm, each theory proposes an understanding of virtue based on the particular exemplar chosen. For MacIntyre, Aristotle's 'virtue' is "a quality which enables an individual to move towards the achievement of the specifically human *telos*" (p. 185). For Homer, it's "a quality which enables an individual to discharge his or her social role" (p. 185). But MacIntyre claims that when these (as well as other) virtue theories are "disentangled," the respective concepts of virtue "*always* require for its application the acceptance of *some prior account of certain features of social and moral life in terms of which it has to be defined and explained*" (p. 186, *emphasis* mine).

In other (my) words, there has to already be in place some macro-contextual arena in which the phenomena of *virtue* is actively exercised and hence provides a framework and backdrop that allows us to reflect on what virtue is. This more fundamental account is that virtue relies on a "particular type of practice" which is

> any coherent and complex form of socially established cooperative human activity through which goods internal to that form of activity are realized in the course of trying to achieve those standards of excellence which are appropriate to, and partially definitive of, that form of activity, with the result that human powers to achieve excellence, and human conceptions of the ends and goods involved, are systematically extended. (p. 187)

One very important aspect of such *practices* is that they require substantial social cooperation and contribution. Consider, for instance, the game of "football" (as opposed to merely "throwing a football with skill"), as

well as the discipline of "architecture" (rather than mere "bricklaying") (p. 187). MacIntyre notes that for the ancients and medievals, the "sustaining of human communities," whether "households, cities, [or] nations," was also this sort of a practice (p. 188). This latter example reflects the definition of a social practice that MacIntyre is interested in, or, what I will hereafter refer to as a *virtue-practice*.[22]

The question we need answered, then, regards what really makes virtue-practices different from all other kinds of practices. Is the difference merely that virtue-practices require more substantial social participation, and hence football and architecture are virtue-practices because *more people* participate in them than the singular practice of 'throwing a football' and 'laying a brick'? Not necessarily. What defines a virtue-practice is not the *amount* of people 'playing the game,' but rather, the fact that it has been originated, shaped, defined, and modified through a substantial "history" of social *participation* (e.g., inter-generationally) that's evidenced in three distinct ways: (1) virtue-practices are handed down (and thus perpetuated), (2) there is a robust social *contribution* to the particular strategies for undertaking and/or modifying virtue-practices, and (3) what can be gained or achieved by undertaking a virtue-practice is socially *articulated* (p. 190). As such, engaging in a virtue-practice not only identifies one's self with the broader community

[22] For the sake of clarity in understanding MacIntyre's argument then, I will hereafter refer to MacIntyrean practices as 'virtue-practices.' Obviously, these are the kinds of practices that he's most interested in.

that also practices it, but at the same time, legitimizes the normative status of that practice as well as that of the macro-context – the intersubjective community of that practice's practitioners – that is involved in the practice (I will discuss this more at length later).

But why do we engage in such practices? For MacIntyre, it is because of the 'internal goods' that are achievable *only* via that practice. Such goods are distinct from the *external* goods that are achievable by undertaking that practice but *not* only by *that* practice. MacIntyre illustrates this with the child who is motivated to win a game of chess in order to win some candy (1981, p. 188). But winning this game isn't the only way to get candy. Furthermore, she doesn't even have to play the game well to get it. She could cheat and still end up with candy. Also, such external goods are scarce. There are only so much of them and that's it. The more one has of them, the less that others can have, and so they become "objects of competition in which there must be losers as well as winners" (p. 190).

What "we may hope" for, however, are not *game-players* who merely want to achieve these kinds of goods, but rather, those *goods* that can only be achieved by playing that particular game – i.e., goods internal to that practice – such as "the achievement of a certain highly particular kind of analytical skill, strategic imagination and competitive intensity, a new set of reasons […] for trying to excel in whatever way the game of chess demands" (p. 188). Here, MacIntyre seems to reflect an affinity to Aristotle's claim that we want (and

need) people to love, and hence develop virtue, because of the good that only that practice can bring. Such love for the practice, then, arises because its internal goods *improve the quality of the person* herself rather than merely increasing the quantity of her possessions. For example, in the case of "real adults," we hope they engage in *social* virtue-practices for the internal goods that contribute to one's improvement as a person rather than for the "prestige, status, and money" they may bring (p. 188). Furthermore, we hope they are experienced at it because as with anything, the ones who are best suited to judge such practices are those who do it often and do it well: "those who lack the relevant experience are incompetent" at judging internal goods particular to that practice (p. 189).

Having laid out an understanding of a 'virtue-practice' in this way, we need ask what is so significant, then, about calling it a 'virtue' practice? For MacIntyre, a 'virtue' is "an acquired human quality" the possession and exercising of which "tends to enable us to achieve those goods which are internal to practices and the lack of which effectively prevents us from achieving any such goods" (1981, p. 191). Virtues, then, are specific and narrow. After all, they are identified in our social practices as those specific tools for achieving the internal goods of a practice and *exclusively* the goods of a particular practice.

The important thing that I think this brings to light on MacIntyre's account is that since virtue-practices are inherently social, and since virtue is

necessary for achieving the goods internal (and exclusive) to particular practices, the argument seems to be that *virtues are the intersubjectively determined norms that govern the practice by which those goods are achieved*. Such goods "can only be achieved by subordinating ourselves within the practice in our relationship to other practitioners" since *a practice is never undertaken solitarily* (p. 191). "Every practice requires a certain kind of relationship between those who participate in it" and "whether we like it or not" virtues are those things "by reference to which […] we define our relationships to those other people with whom *we share* the kind of *purposes* and *standards* which inform practices" (p. 191, *emphasis* mine). It is through these relationships that we gain a clearer understanding as to what are the excellences of the practice that are such *because* they achieve the goods unique to that practice. It would seem that in this sort of social context we regulate the practice itself. As such, virtue-practices are socially rule-governed.

Consider, for instance, MacIntyre's example of the practice of portrait painting. On one hand, the excellences (or *virtues*) of this practice are manifested in the "excellence of [its] products" which in this case are twofold: the "performance by the painter" as well as the excellence "of each portrait itself" (p. 189). Yet on the other hand, there is a sense in which doing the practice itself *is* constitutive (i.e., internal) of a kind of good that is achievable only by doing it. MacIntyre claims that this is "the good of a certain kind of

life" of being a painter and thus "the painter's living out of a greater or lesser part of her life *as a painter*" (p. 190).

But in claiming that this 'kind of life' is a 'good' to be achieved only by this particular virtue-practice presupposes that *others* also are part of the practice in that being identified as a *painter* is to be identified with an intersubjectively governed community that administrates what it is to achieve or be successful in that practice. What this requires "at the very least [is] the kind of competence that is only to be acquired either as a painter or as someone willing to learn systematically what the portrait painter has to teach" (p. 190). So within the community, there are going to be guidelines, standards, rules, and conventions, that shape the practice itself in order to define what it is as well as to define what excellence is within the practice in terms of its external products.

So, "to enter into a practice is to accept the authority of those standards and the inadequacy of my own performance as judged by them" (p. 190). As MacIntyre notes, this is not meant to insinuate that these rules are inflexible or "immune from criticism" (p. 190). Indeed, part of the identifying features of a practice may be how its norms evolve over time, manifesting a legacy of critical engagement that has improved that practice over time. As such, a practice's standards may never be completely constant, but nonetheless "we cannot be initiated into a practice without accepting the authority of the best standards

realized so far" since that is what characterizes the nature of the practice as well as the identity of its practitioners at that time (p. 190).

For MacIntyre, these general facts about our virtue-practices lead to some very specific and interesting conclusions. Because virtue-practices necessarily involve norms or standards, and because those standards are shared, respected, and adhered to through intersubjective discourse of mutual submission to authoritative instruction, "we define our relationship to each other, whether we acknowledge it or not, by reference to standards of truthfulness and trust" (p. 192). For instance, to exercise the excellences of the practice that achieve its internal goods *necessarily* requires one to not cheat (although one could cheat and lie one's way to the external goods) since it would be *impossible* to achieve such internal goods in a way that isn't characteristic of the practice itself. By extension, this also implies that "standards of justice and courage" will be part of any virtue-practice (p. 192). We have to "learn to recognize what is due to whom," "be prepared to take whatever self-endangering risks are demanded along the way," and "listen carefully to what we are told about our own inadequacies and to reply with the same carefulness for the facts" (p. 191). In other words, to *do* the practice excellently, we have to *be* excellent. We have to exhibit those virtues that contribute to being excellent practitioners, for ourselves as well as for others. MacIntyre's example of this is very instructive:

> If A, a professor, gives B and C the grades that their papers deserve, but grades D because he is attracted by D's blue eyes or is repelled by D's dandruff, he has defined his relationship to D differently from his relationship to the other members of the class, whether he wishes it or not. Justice requires that we treat others in respect of merit or desert according to uniform and impersonal standards; to depart from the standards of justice in some particular instance defines our relationship with the relevant person as in some way special or distinctive. (1981, p. 192)

As such, these particular virtues are inevitable since they are inherent to the phenomena of being a practitioner alongside other practitioners. As MacIntyre puts it, "we cannot escape the definition of our relationships in terms of such goods" (p. 192).

But what about the fact that different cultures place different values on different virtues of practices? For instance, "Lutheran pietists brought up their children to believe that they ought to tell the truth to everybody at all times" while "traditional Bantu parents brought up their children not to tell the truth to unknown strangers" (p. 193). How can we say that the virtue of honesty, then, is universally valuable to all practices when it seems as though it actually isn't, as evidenced by the varying codes respective to various cultures?

How Macintyre responds is very instructive. Although codes may differ in degree between different cultures and societies, they at least "embody an acknowledgement of the virtue of truthfulness" (p. 193). Similar practices flourish in different societies despite their different standards, but what they cannot do is "flourish in societies in which the virtues [are] not valued" at all (p. 193). Seemingly, then, the intersubjective nature of engaging in practices –

which involves being critical towards them, modifying them, shaping them, improving them – will necessarily *call for certain standards* or virtues as well as *determine the degree* to which they should obtain.

How do they do this? What is the ultimate standard by which they determine the appropriate norms for a practice? MacIntyre is not so specific in responding to this question. But he does seem to acknowledge the role of a macro-context that governs the standards of virtue-practices. As such, we may speculate how this could be done in light of my arguments made in previous chapters.

Although MacIntyre is concerned with *moral* virtue-practices, it seems that we ought to be able to offer a parallel understanding along the lines of my claim in chapter one regarding epistemic practices. There, I argued that the broader social context determines what the epistemic standards/norms will be for various situations or micro-contexts. How they will vary will be according to *the needs* for that particular community as manifested in the particular epistemic practice it requires.[23] This would explain, for example, why the

[23] One interesting parallel is that virtue, understood as an aspect of a social 'practice' in the sense articulated by MacIntyre, has an interesting affinity with the social epistemic practice discussed by Robert Brandom in terms of his game of 'giving and asking for reasons' that was discussed in chapter one. There, Brandom argued that the practice of this particular game is inescapably a social one where the epistemic community is a normative one in virtue of the fact that it is the authoritative arm defining the standards and procedures for knowledge getting. How we conduct the game, the questions we ask, the conclusions we come to, are not only part of the game but they actually set the standard for the rules that govern that game. As such, epistemology is a sort of game, a sort of MacIntyrean practice, defined by the internal goods (such as justification, for instance, in that it tends to achieve truth and does so in the right way) particular to it as well as by the identity that players of the game take on merely in virtue of being players of the game.

standard for some epistemic virtue may be more stringent in situations where the object of concern is the side effects of a vaccine, and less stringent in situations where the object of concern is whether or not the local 7-11's Slurpee® machine is working properly so that I may determine whether or not my trip there will be a wasted one. Figuring out what our needs are, then, allows us to be clearer on exactly what kinds of goods are most relevant for certain kinds of epistemic practices, such as, as argued in chapter two, lower-grade true beliefs or higher-grade ones – animal and reflective knowledge, respectively. It depends on how we collectively identify and govern the practice.

A question that may naturally arise here, then, is how to regard the role of truth: is it an internal good or an external good? Is it something achievable only by engaging in the epistemic practice of one's macro-context as a 'scorekeeper' or 'gatekeeper' or the like, or can it be attained without being part of any intersubjective epistemic practice whatsoever? One way to think about this may be that the *achievement* of truth is an internal good of our epistemic practice. After all, it's not as though one can 'cheat' at attaining truth (like one could cheat at chess) where such 'cheating' means abdicating oneself from epistemic practices entirely because cheating is merely one's own subversion of the *proper* epistemic behavior specific to a particular epistemic practice in order to attain truth by way of an *improper* sort of epistemic practice. It doesn't seem as though one could remove herself entirely from *rationality*, which seems to

be necessary for recognizing truth, and which we might understand in terms of the plethora of various epistemic practices aimed at the goal of achieving truth that is external to (and hence the reason for the existence of) those practices.

For example, one of my students may cheat on his final exam by refusing to consider what *we* as a class have read and discussed over the course of the entire semester (where our *course* is understood to be a specific epistemic practice among many other intersubjective epistemic practices). Instead, he utilizes the web access feature on his Google Glasses and has full access to the Internet where he finds *true* answers to some of the exam questions. Of course, in a sense, he is obtaining truth *apart from* or *externally to* the epistemic practice of our particular micro-context (i.e., our particular course). Yet, he is still engaging in a truth-achieving epistemic practice. So, it's not clear that one can actually achieve truth through one's abdication from epistemic practices because it's not clear a rational being could ever entirely escape these practices without also escaping rationality itself. In this case, however, the student did not disregard epistemic practices altogether, but chose to engage in an epistemic practice embodied in Wikipedia instead of the one instantiated in *our* course.

Nonetheless, even though he got truth, there is still something amiss here. He is blameworthy for undermining the epistemic practice of our course. Why? After all, he got truth and isn't this the goal? Not exactly. He's not blameworthy for merely obtaining truth, but for doing so in an illegitimate

way given the purposes, interests, needs, and rules of the *specific* practice of our course. In other words, the student obtained truth in an unjustifiable way according to the norms and rules of our epistemic practice. What is valued most here (and rewarded at the end of term), then, is not merely possessing truth, but having achieved it in the right (*justified*) way where what it means to 'be justified' is determined by the rules of the practice itself, which are influenced by the purposes, interests, and needs that give life to the practice. As such, *justification* itself seems to characterize epistemic practices since such practices are characterized by what they do: they evaluate, criticize, appraise, and either accept (i.e., *justify*) or reject our belief claims.

One potential issue that may arise from this discussion, however, is that if truth is the goal of our epistemic practices, wouldn't this mean that it is defined by our practices in the way that goals such as 'touchdowns' and 'checkmates' are defined by the practice of their respective games? After all, the possession of truth (i.e., knowledge) is the goal of the epistemic game as touchdowns and checkmates are the goals of football and chess, and it turns out that these latter goals have been created and defined *out of* the practice of those games. So why wouldn't 'truth' be the same, that is, just a creation of our epistemic practices?

One way to answer this is that while truth itself may *not* be *defined* by our justificatory practices but, at least in part, by how the world is, (and may even in some cases, for example, regarding animals' beliefs about their

environment, be achieved apart from participation in justificatory practices), *the proper and appropriate way* for us and rational beings like us to obtain truth is indeed *defined* by those practices.[24] As such, it's the recognition of and motivation for truth that brings about the creation, refinement, and even the variety of our epistemic practices which are geared toward achieving the truth *in a certain kind of way*. This is why my cheating student is blameworthy – not because he doesn't possess truth (because he does, assuming Wikipedia is correct), but because he (intentionally and maliciously) abdicated his responsibilities as an epistemic practitioner in our particular practice which means that he didn't obtain the truth in a justifiable way (given how our practice understands justification). As a result, this student doesn't enjoy the goods internal to our practice – possessing truth *because* he is justified (i.e., has appropriately/properly obtained those beliefs in a way that is sanctioned by our practice). Hence, while *truth* is a good that is external to our specific practices, *obtaining truth in our practice-specific (i.e., justified) way* is indeed a good that's

[24] What I have in mind here is Ernest Sosa's (1991) criticism of Pollock's criteriology according to which our concepts, including truth, are *constituted by the rules* or norms or criteria that govern their application in our cognitive practices. And, so, there is no such thing as good reasoning with our concepts, including truth, understandable apart from the norms of our actual epistemic practice, and thus no deeper unity to good epistemic practices such as critically seeking and obtaining truth. However, as Sosa notes, "the unity of chess is that of an invented practice, invented and adopted because it coheres nicely and defines a challenging and amusing activity" and because of this, "it is hard to see reason as conventionally invented" or as "constituted by a set of rules held together by nothing more than their mutual adjustment for the sake of an amusing challenge to humans" (p. 130). As such, Sosa seems to suggest here that our epistemic concepts like the goal of 'truth,' for instance, has some sense of being and meaning apart from any practice we refine and modify for the purpose of achieving it. Hence, our epistemic practices get their being and purpose from our recognition of their ability and propensity to get us truth.

internal to our epistemic practices since one must exclusively follow its norms and procedures in order to achieve this good.

Finally, one might speculate that a certain parallel may be found between this discussion of Macintyre and Aristotle's argument that a theory of virtue can only be given in outline. On the one hand, there are those basic necessary excellences of any practice (e.g., Aristotle's list of virtues for the practice of achieving the 'good' within his particular social context). On the other hand, however, he (as I argue, anyway) seems to recognizes that features of the micro-context itself is going to determine how a virtue should be particularly manifested, which is to say that there is a sense in which the social governance aspect of the ethical practice (i.e., one's relationships and attitudes and perceptions at the macro level) will determine what the norms/rules are for virtue in particular situations. As such, it's not as though the normative role of the community is undermined through varying degrees of standards and rules of the practice, but rather that the standards for identifying and refining the specific appropriate manifestation of virtue for particular micro-contexts may evolve with the practice itself in order to fit the practice better (and make the practice a more excellent one). This reaffirms the normative role of the macro-context.

As MacIntyre notes, "practices never have a goal or goals fixed for all time […] but the goals themselves are transmuted by the history of the activity" (1981, pp. 193-4). As such, "every practice has its own history" and

this "historical dimension is crucial in relation to the virtues" since it is the story of the relationship we have with previous practitioners of the practice (p. 194). As an intersubjective and intergenerational collective, then, we take the shape of the staggered non-uniform strands that form continuity within the practice itself (p. 194). As ManIntyre notes, especially regarding the practice of seeking human good, none of us are ever able to "seek the good or exercise the virtues" as mere individuals since "we all approach our own circumstances as bearers of a particular social identity" (p. 220). As such, it's not really the case that we all 'choose who we want to be,' but rather, we inherit conceptions of who we are in virtue of the history of the practices we engage in. As a result, "I find myself part of a history and that is generally to say, whether I like it or not, whether I recognize it or not, one of the bearers of a tradition" since "virtues sustain the relationships required for practices" no matter of those relationships are in "the past – and to the future – as well as in the present" (p. 221).

3.5. If Virtue is Contextual, Does That Mean it's Not Real?

A challenge that this contextualist understanding of Aristotelian-styled virtue theories may raise is the notion that such an ethical framework is, in fact, an artificial one rather than an organic one, so to say, that it arises as a result of being humanly constructed rather than discovered. After all, I suggested that for Aristotle, the *recognition* of 'moral truth' (i.e., moral

judgments, virtues, etc.) is the work of the 'well brought up' folks, as noted earlier in this chapter, those who are of the superior statuses – i.e., the ruling/wise classes. Why is this work of recognition *their* burden? Because, after all, they possess the *phronësis* (practical reason) and *synesis* (understanding) that their particular lifestyle allows them to develop due to the privilege their pedigree affords them. And so, I might be viewed as suggesting that when they determine what the virtues are, they are really *determining* them in a strong *creative* sense.

The challenge, then, may go something like this: the way that I have presented Aristotle and the tradition that follows him, then, suggests that for him virtue doesn't really exist *per se*, so to say, but *per accidens* – not essentially on its own, but contingent upon being created (through judgment) and agreed upon by the macro-context (i.e., the well brought up folks). Effectively, then, virtue is *determined* via the macro-contextual act of praising and prescribing it in the model of the exemplar, and hence is not something that is *discovered* and deemed praiseworthy (and then desired) because of its optimal contribution to human function and *eudaimonia*.

If this is correct, and the rightness of a moral judgment (virtue) is not merely recognized as such but is actually created by the macro-context (and it is in virtue of its act of creation that confers moral rightness upon it), then what may actually be suggested thus far is that Aristotle's ethical framework may not merely be interpreted in a contextualist light, but that he is quite the

relativist in a most extreme way. In short, then, the question is this: am I inevitably arguing that Aristotle is a non-realist regarding moral correctness, that there doesn't exist any non-perspectival fact regarding what is the right thing to do in a situation?

Given this complaint, let's break it down into two claims. Am I arguing in this chapter that (1) this is Aristotle's view, that the judgment of the virtuous exemplar is what *constitutes* rightness, and hence the metaphysical status of the *correct* moral judgment is not ontologically independent of how the macro-context understands it to be?[25] And furthermore, am I arguing that (2) for Aristotle, the community of virtuous persons' judgment as to what are the standards for *manifesting* a virtue define/constitute/make-for the *correctness* of those standards?

There are good reasons to reject claim (1), and I in no way intend to defend or argue for it.

Consider Terence Irwin's (1988) famous review of Martha Nussbaum's *The Fragility of Goodness* (1986) where he criticizes her[26] for arguing that Aristotle's moral theory exhibits a rejection of realism when it's not clear

[25] This seems to be Martin Tweedale's position: "There is no fact of the matter beyond what the discussants in an ethical inquiry agree on from the start or what a rational discussion among these people suited to ethical inquiry can convince them of. To put it bluntly, Aristotle is not a 'realist' in our sense when it comes to truth in matters of ethics, and that fits perfectly well with his pragmatic conception of ethical inquiry not as aimed at achieving knowledge but rather at improving ourselves" (2007, p. 6).

[26] "Nussbaum's treatment of the texts she discusses displays a pattern of distortion and omission whose cumulative effect is to conceal evidence that appears to conflict with her general view. Readers will be well advised to check her translations and paraphrases of Greek texts and her descriptions of the views of modern philosophers and scholars. The faults in the argument of this book justify skepticism about its main claims" (Irwin, p. 383).

whether Aristotle does or not. Irwin notes that while Nussbaum tries to argue that Plato is a metaphysical realist regarding the wrongness of certain kinds of actions independent of "any human agreement," Aristotle, on the other hand, "accepts some human agreements as having evidential value toward his conclusions" and thus "Aristotle is not a metaphysical realist about his conclusions" (Irwin 1988, p. 378). For Irwin, Nussbaum's conclusion about Aristotle would make more sense if she could attribute to Aristotle the view that his epistemological position regarding finding moral truth is "nonperspectival and nonanthropocentric" – i.e. a 'view from nowhere – but she doesn't do this (p. 378). For Irwin, however, there seems to be at root a fundamental mistake that Nussbaum makes which forces her misinterpretation of what's going on in Plato and Aristotle:

> She seems to rely on a confusion about relativity. What is healthy for a human being depends on a human being's nature, including human imperfections. But a more perfect being can recognize these facts, and can therefore see what is really healthy for a human being. Taking up a 'nonanthropocentric' or 'god's-eye' or 'nonperspectival' point of view does not mean that we cannot recognize values that are valuable because of imperfections (Irwin 1988, p. 379).

What I take from this criticism of Nussbaum is that whether or not Aristotle indeed has some sort of 'view from nowhere' regarding what makes for moral virtue, either way, it doesn't necessarily commit him to moral non-realism. After all, the more important point for Aristotle is that human beings are morally and cognitively structured a certain way – what he identifies as their 'function.' So, given that they're structured so as to function a certain way

(which Aristotle seems to accept as a fact about humans), it would seem that there would naturally arise additional facts about the *correct* way for such human beings to function to the highest and most optimal degree possible, that is, according to that function.

Put a different way, *relative* to the way in which human beings actually *are*, there must inevitably also exist facts of the matter regarding how those beings can best live up to the way in which they are, and for Aristotle, such a life is going to be one that evidences the practice of virtue. So, what are going to be moral virtues will be *dependent* on human beings in one sense (they are non-intelligible apart from the particular way that human beings are structured), but *independent* of human beings in another sense (given that human beings are structured the way that they are, what are the correct moral virtues will inevitably be so, regardless of whether or not they are recognized as such by those beings, even though it would seem that the 'bright' 'well-raised' folks will figure this out).

Seemingly, then, this may be why Aristotle argues that *relative* to the way human beings *are*, it is *factually* the case that actions such as "adultery, theft, and murder" and passions such as "spite, shamelessness, [and] envy" are unequivocally bad for human beings (1107a10). And, also, relative to the way in which human beings actually *are* (i.e., the way the world actually is) it is factually the case that virtues such as 'generosity' and 'courage,' for instance, are unequivocally good for human beings. But there is a further complex aspect

of these latter kinds of virtues in that it seems unlikely that they could be appropriately manifested the *same* way, to the *same* degree, in all of the different situations that calls for their exercise. Hence, Aristotle argues for the unequivocal moral wrongness of 'adultery' as well as he argues for the unequivocal moral rightness of 'courage.' But courage is quite different, structurally so, since its appropriate manifestation depends on contextual features of the situation that calls for its exercise.

Regarding claim (2), then, it does seem to be the case that the standards for determining how a moral virtue is to be appropriately (correctly) *manifested* in a particular situation will depend on the norms given for that particular situation (or micro-context) by the determinations of the macro-context. As such, what the macro-context really does is play an important epistemological role through the model of the exemplar as that which expresses what are the moral norms. The exemplar is set up by the macro-context as a standard that the community follows. So, in a sense, we might say that the standard (exemplar) is ontologically dependent on the community. However, what makes that standard (exemplar) a *good* one depends on human nature in terms of what our human function is, as well as the demands of the various micro-contexts themselves. And, this is *not* the same thing as claiming that the macro-context *creates* virtue. Rather, it is merely to say that the macro-context discovers the moral excellences as part of its exercise in unearthing the way the world (i.e., rational beings) actually is and how it

flourishes (i.e., pursuing *eudaimonia*) in light of what it is. As such, it seems that for Aristotle's ethical framework, the macro-context may really just serve an important epistemological and pedagogical role as a sort of 'epistemic conduit' for the dissemination of moral norms and standards whereby we judge and know what the right thing to do is.[27]

An argument along these lines also might apply to MacIntyre's conception of virtue-practices. As Aristotle seems to suggest that virtue is an inevitable fact that arises out of our particular function as rational beings, MacIntyre seems to argue that virtue inevitably arises out of any kind of practice that is geared toward achieving the internal goods of that practice. After all, for MacIntyre, virtue-practices necessarily involve standards and norms – such practices, then, will inevitably reveal virtues. For instance, MacIntyre noted that within a practice its participants "define [their] relationship to each other, whether [they] acknowledge it or not, by reference to standards of truthfulness and trust" as well as by "standards of justice and courage" (MacIntyre 1982, p. 192). Why? Because, seemingly, any practice that is intersubjective (and hence is a shared practice) will necessarily rely on

[27] This understanding of the epistemic role of the macro-context regarding *moral* norms, however, seems to diverge from my chapter-one argument regarding the role of the macro-context concerning *epistemic* norms where 'scorekeepers' and 'gatekeepers' actually determine (i.e., create) the truth-conditions for knowledge. Of course, one thing that both kinds of macro-contextual understandings seem to have in common, however, is that the purposes, needs, and interests of the community influence their respective resulting standards/norms. Whereas on the Aristotelian picture 'needs, purposes, and interests' will be that we flourish (in the various micro-contexts) in terms of our human function and human nature, on the score-/gate-keeping model, 'needs, purposes, and interests' have to do with how serious the 'stakes' are (in various micro-contexts) and hence what epistemic judgment will be in our best interest. As such, it may be that moral and epistemic norms are a response to these 'needs, purposes, and interests.'

these for its success and integrity. That's just a fact about the way practices work – just a *fact* about the way the world is relative to that human phenomenon.

As MacIntyre goes on to say, we have to "learn to recognize what is due to whom," "be prepared to take whatever self-endangering risks are demanded along the way," and "listen carefully to what we are told about our own inadequacies and to reply with the same carefulness for the facts" (p. 191). The need for such virtues, then, reveals that they are part and parcel of the practice and that engaging in the practice is how they are discovered. But this isn't to suggest that the community of practitioners creates the virtue upon recognizing it. Arguably, they would discover its necessity since it is inevitability tied to the practice itself in a factual way, that is, as relative to the notion of a virtue practice. But, then, how that gets *particularly* cashed out or manifested in one sort of practice (such as football) or in another (such as architecture), will be *relative* to the intersubjectively *determined* parameters and norms for that practice by those who are authoritative within that practice.

As such, it's not clear that MacIntyre would be as objective about virtue as it seems that Aristotle is. Although, as just noted in the previous paragraphs, MacIntyre argues that there are inherent universal virtues (like honesty and justice) that automatically arise via the nature of practices themselves, it also seems that he considers the macro-context playing a much more extensive (i.e., ontologically creative) role since the practitioners within a

practice will influence its trajectory in the sense of determining what it will mean for a practice to be done well. After all, MacIntyre doesn't seem to share Aristotle's 'perfectionist teleology' where, given the way that human beings are (i.e., structured/functioning in a particular way that determines what flourishing is), certain kinds of moral virtues inevitably result. MacIntyre's view, then, doesn't depend on Aristotle's metaphysics of human nature in terms of function, purpose, and the flourishing respective to such function and purpose. As such, there may be a bit more room for relativism on MacIntrye's account than on Aristotle's.

To sum up, then, what I have tried to offer in this section is a neo-Arsitotelian interpretation that considers virtues to be *relative* in two specific ways.

First, they are relative to the way in which rational beings are structured and this is just a *fact* about humanity. Macro-contextually, then, this is recognized through *discovery* (due to the exercise of *phronësis* and *synesis*) of who we are and what is best for us. And hence, once these discoveries are linked, a particular (appropriate) manifestation of virtue is prescribed.

But, secondly, the best and most appropriate way to apply the virtues (i.e., the most appropriate way to manifest them in particular situations) in specific micro-contexts will be relative to that micro-context – i.e., to the cultural expectations, norms, needs, interests, purposes, etc., recognized by the macro-context and prescribed in the form of norms and rules (and an

exemplar) for the appropriate manifestation of virtues in those particular micro-contexts. Hence, where a certain sort of 'perspectivism,' if you will, is appropriate (and even necessary) is in how to correctly manifest that virtue – and it seems that that is going to be culturally (i.e. macro-contextually) relative. After all, for example, even though courage may be identified as a virtue given how human beings actually are, how it should be manifested on a battlefield as opposed to during one's final exam will be very different.

3.6. Conclusion: Transitioning to Responsibilist Epistemology

What I have been running in this chapter, then, is my own neo-Aristotelian position regarding virtue and the contextual variation of its manifestation. What I have tried to show, then, with evidence from the work of Simpson, Tweedale, and MacIntyre, is the indispensible role of the intersubjective community in discovering and applying virtue. MacIntyre specifically offers a conception of virtue theory where the macro-contextual aspect of virtue is much broader than what Aristotle explicitly argued. As such, MacIntyre's theory is a good example of how virtue can be macro-contextual and yet not necessarily rely on a specific understanding of a moral exemplar. Rather, emphasis is on the particular community of practitioners.

One claim might be, however, that MacIntyre may not be too far off from Aristotle (or, that in some way maybe they are harmonizable) since the notion of a moral exemplar might possibly be the kind of thing that a

community comes up with in order to serve as the standard-bearer for its practices. What will be interesting, then, is to see how this might apply to an epistemological theory that is also neo-Aristotelian, like the virtue responsibilism that Linda Zagzebski offers, and to which we turn next.

Chapter Four

Virtue and Context in Responsibilist Epistemology

> *The problem here is that virtuous procedures underdetermine action and belief in the particular case even when there is no conflict among the relevant virtues and even when finding the mean is not the salient issue. […] Even when we list all of the known procedures used by a virtuous person, there is still an indeterminacy with respect to most beliefs. […] Philosophers often make intellectual caution a matter of professional pride, but in my experience, the beliefs of philosophers on philosophical issues exceed that which they can justify to others or even to themselves on the basis of following commonly accepted epistemic procedures. […] Phronësis is needed, then, [because] many human activities, whether of the overt kind traditionally handled by ethics or the internal activities of thinking and forming beliefs, can be neither fully described or evaluated in terms of the following of a set of known procedures or rules. Good judgment is required in all areas of human activity, including the cognitive.*
>
> – Linda Zagzebski, 1996

4.1. Moving From Ethics to Epistemology

In the previous chapter, we considered a possible contextual reinterpretation of Aristotelian basic ethical framework as well as Alasdair MacIntyre's notion of virtue as arising out of our intersubjective social practices. In doing so, we got a very clear glimpse of how the macro-context and micro-context works on these accounts as well as how they are connected. Given that much of the arguments for responsibilist epistemology are based on

the broader virtue ethical tradition, we looked at some examples of the latter in the previous chapter to set some groundwork for this one where responsibilist virtue epistemology is the topic.

I intend to argue that such an epistemology mirrors the ethics it's based upon, not merely in terms virtues which achieve their target 'goods' (here, understood as correct beliefs), but that, in some sense, they are determined by the community, the macro-context, in such a way that then determines how that should be played out in various particular epistemic situations. In what follows, I consider two virtue-responsibilist pioneers – Lorraine Code and Linda Zagzebski – and argue that implicit in their work is some sense of a macro and micro contextual aspect as I have understood it to be thus far.

4.2. Motivating a Contextual Responsibilist Epistemology

Lorraine Code, after finding inspiration and motivation from Ernest Sosa's first attempt at articulating a theory of justification based on the notion of *reliabilist* abilities or competences,[1] wrote her magnum opus, *Epistemic Responsibility* (1987), where she offered the inklings of a broad framework for a possible virtue-epistemology based on a *responsibilist* understanding of virtue, that is, one based in the more traditional Aristotelian sense rather than Sosa's

[1] Ernest Sosa (1980). "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3-25.

sense.[2] "Disenchanted" by the restrictive elements of foundationalist and coherentist justification due to their guiding vision of knowledge as an "individual" attainment, Code discusses how the "knowledge seeking activity" of human agents is both a communal and moral practice (pp. 7-8). For her, "responsible knowers are uneasy about possible inconsistency and incoherence" not only "within their own system of knowledge" but also "within 'public' systems of knowledge" (p. 6).

As a way of illustrating her concerns, Code considers the intellectual life of scientist Sir Edmund Gosse, the subject of the 1907 memoir *Father and Son: a Study of Two Temperaments*, authored by his son Philip Gosse. This memoir discusses Edmund's intellectual crisis of being faced with the tough epistemic choice of either adhering to a creationist belief-system (rooted in his fundamentalist religious leanings) or be open to the new evidence (and its alternative explanatory implications) provided by Darwinian theory which was becoming increasingly accepted by the wider scientific community. As such, this problem involved "the boundaries of two ultimately conflicting systems" (1987, p. 21). But in the end, Gosse chose to maintain his current beliefs instead of opening them up to the scrutiny provided by Darwinianism. As Code notes, Gosse was "not prepared to consider that it might be possible, and

[2] As Code notes, "I call my position 'responsibilism' in contradistinction to Sosa's proposed 'reliabilism', at least when human knowledge is under discussion […] because the concept 'responsibility' can allow emphasis upon the active nature of knowers/believers, whereas the concept 'reliability' cannot" (1987, p. 51).

indeed desirable, to look at one's putative knowledge from the outside" since "he seems to have locked himself into a mode of intellectual practice that does not […] allow a self-critical stance" (p. 21).

For Code, Gosse's failure is one of "reflexivity" in that he is "quite unaware of his own dogmatism" that leads "ultimately to a failure in integrity, wisdom, and epistemic responsibility" (p. 23). In a word, then, Gosse's failure is a *failure of virtue*, which for Code, is a failure that precludes a putative knower from being an accurate knower.

> Intellectual virtue is, above all, a matter of orientation toward the world, toward one's knowledge-seeking self, and toward other such selves as part of the world. Central to it is a sort of openness to how things are: a respect for the normative force of 'realism.' This attitude involves a willingness to let things speak for themselves, a kind of humility towards the experienced world that curbs any desire to impose one's cognitive structurings upon it. Intellectual honesty consists in a finely tuned balancing of these two factors, in cultivating an appropriate interplay between the self and world (1987, p. 20).

For Code, an approach to epistemology based on responsibilist virtues such as wisdom (p. 53), intelligence (p. 55), and prudence (p. 55), among others, is ideal in that it helps to explain three factors that she argues are pivotal for understanding the nature and targets of our epistemic practices.

First, "knowledge claims and efforts to know are events or processes […] that emerge out of the interaction amongst knowledge seekers, their communities, and the world" (p. 26). As such, knowledge-norms arise out of the needs and practices of macro communities whether such a community is local (narrow) or global (broad). After all, the norms that arise out of our

practices suggest that we don't merely want to be successful knowers but we want others to be as well. This is further manifested in what Code argues is a fact about being intellectually virtuous: it's not just a matter of having one's epistemic endeavors "come out right," but that one possesses "a fairly constant and dependable set of qualities and capacities" that are not only manifested in one's "orientation toward the world" and "toward one's knowledge-seeking self," but also "towards other such selves as part of the world" (p. 52). As such, intellectual *virtue* becomes the standard by which we gauge the epistemic integrity of each other (in terms of whether or not we possess/exercise such virtue) and thus the norm by which be either attribute knowledge to those who are praiseworthy, or withhold such an attribution from those who are blameworthy (i.e. such as the elder Gosse).

Second, as a result of this, "there is no knowledge without knowers, no knowledge without context" (p. 26). Here, Code seems to have in mind the *macro*-context, as I have termed it. To speak of 'knowledge' solely as an individual practice is quite inaccurate, not only because we need others for the terms of our language (such as 'knowledge') to make sense, but because it is a communal and social practice carrying standards and norms for its practice that must be determined and communicated so that knowledge can be disseminated and attributed.

And, finally, "knowledge cannot be stored equally in a computer or a human mind because people have attitudes to knowledge that shape both its

structure and its context" (p. 26). What this suggests is that knowledge is not the mere possession of data. It's not the grasping of facts about the world. Rather, it's the possession of those facts or data (or true propositions) attained in a particular way, in accordance with certain norms that sheds light on our attitudes about the objects of our knowledge practices as well as how we value it. For example, consider my true beliefs that

($G^{13,457}$) There are 13,457 blades of grass on my front lawn.
(GS^{AB}) I have been diagnosed with Gilbert's Syndrome.[3]

Why is it the case that my having the true belief GS^{AB} is more important and more valuable to me (but not only to me) than the true belief about the number of blades of grass on my front lawn? Well, there are many reasons. One may be, to use the language of chapter one, that there are 'higher stakes' involved with whether or not GS^{AB} is true or false than if $G^{13,457}$ is true or false in regards to my health and well-being. So, if I care about my own health (and if you do too), then we both care more about the truth-value of my belief GS^{AB}. But even if you don't care particularly about my health, you still might care about the truth-value of that belief since the standards which govern whether or not it is a justified one for me would also govern whether or not it is a justified one for you, and thus share an affinity to *your* beliefs about your own health and of others that you do care about. As such, for many reasons,

[3] Gilbert's Syndrome is a genetically inherited liver condition that causes a person's bilirubin level to run higher than what is considered normal, thus causing the body to run at a higher than usual level of toxicity, but usually not so high as to cause significant damage.

we agree that our attitudes about this kind of belief (regarding getting diagnoses correct and our process for doing so) does indeed affect how we go about forming and scrutinizing beliefs of this kind. Such attitudes, then, influence what our standards are for getting true beliefs, attributing them, and defining knowledge.

Code, then, is explicit in that our epistemic practices are inherently macro-contextual since our knowledge claims and conception of 'knowers' are intelligible only because they take part within a larger epistemic community. Attributions of knowledge are shaped in virtue of who *we* are as knowers, who in turn are epistemically shaped and conditioned in virtue of being part of that epistemic community. The responsibilist approach to epistemology, then,

> denies the autonomy of the known, maintaining that the nature of the knower and of his/her environment and epistemic community are epistemically relevant, for they act as enabling and/or constraining factors in the growth of knowledge, both for individuals and for communities. […] It does seem true […] that one cannot hope to understand human action in isolation from lives, histories, contexts, and narratives, and I think it is equally true that one cannot hope to understand cognitive activity and intellectual virtue apart from lives, histories, and contexts (1987, p. 27-8).

Of course, for Code, understanding epistemic *responsibility* in terms of whether or not knowledge attributions exhibit intellectual virtue suggests a strong ethical component to epistemic activity since the "knower," and not merely "the known," is "accorded epistemological significance" (1987, p. 37). After all, when we talk about ethics and humans' role in it, it is the *agent* that is responsible for acting ethically and is thus attributed as being (or not being)

such, and not merely her actions as if they are divorced from her responsibility (which they aren't). And since Code thinks she has successfully argued that epistemic failures are failures of virtue, that is, failure of living up to the socially determined standard for what counts as knowledge-conducive epistemic behavior, it is not the agent's false belief itself that is responsible for her failure, but the failure of her *self* which also turns out to be an instance where she 'fails the community' too since she disregarded the norms and acted epistemically irresponsible.

For Code, in this way, "viewing epistemology as analogous to ethics provides a useful perspective on epistemological questions" because it gives us an automatic sense of how to understand the norms or rules that will govern the practice: virtues (p. 38). As such, "the center of epistemic focus" in terms of those norms or rules for belief justification "will become an intellectual analogue of the stable virtues and dispositions" because having an epistemology "modeled on this paradigm" means that "evaluative significance will be accorded to [the subject's] disposition" (p. 43-4). As such, epistemic integrity will not only be a *result* of virtuous intellectual dispositions but will also influence how such dispositions ought to look when re-conceptualizing and modifying the framework of norms for epistemic justification. From this, then, we get the criteria for what counts as good epistemic behavior:

> This endeavor will, of necessity, appeal to various criteria of virtuous cognitive conduct. In other words, epistemological judgment is not just a matter of

> assessing individual conduct *per se*, but of assessing it as a manifestation of justifiable social practices and approaches to enquiry (1987, p. 44).

With this we then get social practices of epistemic inquiry that will require of people, for particular kinds of performances in particular micro-contexts, "certain roles and sets of circumstances [that] impose standards of intellectual achievement over and above those expected of persons simply as persons" (Code 1987, p. 62). In other words, there is a macro-contextual expectation that certain individuals, given their epistemic skill-sets for certain kinds of micro-contexts, will be epistemically responsible in those micro-contexts for the sake of what belief-claims get disseminated throughout the community as well as for the sake of how the appraisal of those particular claims will influence further epistemic practices and standards.

Referring to Gosse again, Code notes that because he is a "respected scientist and a Fellow of the Royal Society, [he] has epistemic responsibilities that, in a sense, transcend those of 'ordinary' members of an epistemic community" (p. 62). Gosse and others like him, then, face "demands of epistemic responsibility of a more pressing nature than those that face an 'average' enquirer" (p. 62). In particular, Code claims that this is so because he has a significant role within the macro-context: "he is one of those who shape the standards of responsible enquiry" that all the rest of us abide by (p. 62).[4]

[4] One interesting parallel here is how this fits with my earlier argument, as noted in chapter one, regarding the 'stakes' involved with a particular person's knowledge claims given his particular task in the broader macro-context. As Sarah Wright has noted, this points to the "relevance of the social role of the potential knower" in that such "relevant social roles are a

But, also, given that "human beings are cognitively interdependent in a fundamental sense, we would expect him to have intellectual integrity, not only in his professional life, but in his personal life as well" (p. 167).

In other words, from him and others like him, we have "expectations of integrity" such that "are closely connected with the attributions of intellectual virtue" (p. 63). As such, *we* want our practitioners of knowledge to have intellectual virtue since such virtue will guide their epistemic practice and knowledge claims: "intellectual virtue manifests itself in communities that impose constraints and conditions upon acceptable knowledge seeking at the same time as they make that activity possible" (p. 253).

4.3. Zagzebski's Aristotelian-Based Responsibilist Epistemology

In an attempt to articulate this sort of epistemological approach more thoroughly than Code, Linda Zagzebski offers a theory of responsibilist epistemology that is based on the framework of Aristotle's moral theory.[5] In

---

feature of the subject's and not the attributor's context," that is, the particular role itself is an aspect of the micro-context in that it affects what will be the epistemic behavior and performance we would expect from someone filling that role in that situation and thus, as Code would probably say, picks out their particular epistemic responsibility in it (Wright 2011, pp. 103-4).

[5] An important thing to note here is that while Zagzebski ultimately argues that intellectual virtues are a "subset of the moral virtues in the Aristotelian sense of the latter" (1996, p. 139), Code claims that one needn't have to do think of virtues in such a narrow way in order to have a responsibilist epistemology: "it is perfectly reasonable to argue […] that epistemic concepts are not reducible to ethical concepts; it is quite another matter to propose, as I am doing, that we structure our epistemological reasoning on an analogy with our moral reasoning […] this does not amount to an insistence that we separate the moral from the epistemic uses of terms such as 'right, wrong, good, bad, justified, unjustified.' The point is to understand the similarities and differences in the reasoning processes that warrant the application of these terms" (Code 1987, p. 48). Zagzebski, however, seems to do both things

the prefatory comments to her 1996 work, she notes her dissatisfaction with the plethora of virtue epistemologies previously articulated because none truly present what she thinks is a properly aretaic account. For her, if a theory claims to be one of virtue, it ought to truly reflect an *aretaic* structure, such as Aristotle's, rather than merely stipulating a very general notion of virtue as a way to talk about a mere faculty/ability/skill. As such, she criticizes how Ernest Sosa (and other reliabilists) talk about virtue as merely a 'competence' or 'ability' without including other important aspects of traditional virtue theory such as a virtue's motivational capacity and its relationship to one's moral/intellectual character.[6]

---

here: she wants to argue that epistemic virtue is analogous to moral virtue and thus that our epistemic reasoning is analogous to moral reasoning, while also arguing that the epistemic enterprise is an aspect of the broader moral one.

[6] As Zagzebski chides, "[Sosa] makes no attempt to integrate intellectual virtue into the broader context of a subject's psychic structure in the way that has been done by many philosophers for the moral virtues. […] His plea for a return to the concept of intellectual virtue actually has little to do with the concepts of intellectual virtue as virtue in the classical sense" (Zagzebski 1996, pp. 8-9). Sosa, however, thinks that his notion, although not "a narrow Aristotelian conception of virtue according to which a virtue is a certain disposition to make appropriate deliberate choices," is plausibly a "broader sense of 'virtue,' still Greek, in which anything with a function – natural or artificial – does have virtues" (Sosa 1991, p. 271). Zagzebski concedes that "in one sense of virtue, Sosa [is] right" but she wants to think of virtue not merely as a "properly functioning natural cognitive process," but "the narrower concept" which understands it as "an acquired human excellence […] for which we are responsible" (1996, p. 103). Both Sosa and Zagzebski, then, define 'intellectual virtue' in their respective ways with some basis in Aristotle, but neither really carries over Aristotle's definition. For him, the intellectual virtues (wisdom, scientific knowledge, theoretical reason, practical reason, and art) were the excellences of the rational aspect of human beings that were necessary for correct or truthful thinking regarding the kinds of objects of knowledge that the particular intellectual virtue would be suited for. The one most relevant to ethics is the intellectual virtue of *phronësis* or practical reason which he argued was necessary (as we saw in the previous chapter) for not only having knowledge of the moral virtues but more importantly in figuring out how to appropriately manifest them in various situation, i.e. finding the 'intermediate.' Zagzebski, however, subsumes intellectual virtue into the moral virtues by arguing that "an intellectual virtue does not differ from certain moral virtues any more than one moral virtue differs from another" and that "the processes related to the two

In contradistinction to this, Zagzebski takes an Aristotelian tack and argues that intellectual (or epistemic) virtues aren't merely skills or abilities, but are parallels to the moral virtues: an intellectual virtue is "a deep quality of a person" (p. 104) with a "motivational component" (p. 115) that acts "within us" in that it "initiates and directs action toward an end" (p. 129) with "persistent tendency" (p. 132). For Zagzebski, possessing the motivation to believe true propositions "leads a person to follow rules and belief-forming processes that are truth-conducive," which ultimately means that agents who are virtuously motivated are motivated to develop and exercise the particular intellectual virtues and their subservient skills that are conducive to (and necessary for) achieving true beliefs (p. 176).

Let's examine the particulars of this epistemological approach.

First, consider Zagzebski's notion of what a *virtue* is. Her project emphasizes the Aristotelian maxim that *virtues are excellences*, understood as "deep traits" of its possessor that mirror the character and actions of the virtuous exemplar (1996, p. 100). That virtue is a "state of the soul" insinuates that it is a property we attribute to a person "in a deep and important sense" that is "closely associated with her identity" (p. 85). Just as with Aristotle, for Zagzebski, a virtue makes its possessor good in the sense that it increases her overall moral worth (p. 95).

---

kinds of virtue do not function independently," hence, "intellectual virtues are best viewed as forms of moral virtue" (Zagzebski 1996, p. 139).

But Zagzebski thinks that such descriptions are true of *intellectual* virtues as well as moral ones (p. 95). A person's "intellectual worth" is greater for the person who possesses intellectual virtues such as open-mindedness, intellectual carefulness, being "thorough in evaluating evidence," being "intellectually courageous," and so on (p. 95). Intellectual virtues like these are not innate abilities that are part of our normal function or design, but are "acquired human excellences" for which we hold others and ourselves responsible (1996, pp. 100-1). Following Code, Zagzebski argues that this is something we observe from our practical dealings in the world. We don't hold people responsible for their 'natural' capacities, talents, or even looks, but we do hold them responsible in matters where they have a choice. We consider them praiseworthy when they possess virtues and blameworthy if they lack them (pp. 104). Consider Zagzebski's working list of intellectual virtues:

> - The ability to recognize the salient facts; sensitivity to detail;
> - Open-mindedness in collecting and appraising evidence;
> - Fairness in evaluating the arguments of others;
> - Intellectual humility
> - Intellectual perseverance, diligence, care, and thoroughness;
> - Adaptability of intellect;
> - The detective's virtues: thinking of coherent explanations of the facts;
> - Being able to recognize reliable authority;
> - Insight into persons, problems, theories;
> - The teaching virtues: the social virtues of being communicative, including intellectual candor and knowing your audience and how they respond.
> - Originality, creativity, inventiveness. (1996, pp. 114, 182)[7]

[7] It's worth mentioning that if any of these were to conflict in a particular circumstance, it's thought that *phronësis* would settle the matter.

Second, consider the role for 'motivation' in Zagzebski's account. For Aristotle, one becomes motivated to become good once one has had a 'taste' of the inherent goodness of virtue. Such a taste incites the agent to come to love virtue and then want to cultivate it even more by developing those character traits that are the virtues themselves. In this way, experiencing the inherent goodness of virtue itself (given its conduciveness to *eudaimonia*) makes the agent want more of it. For Zagzebski, this "characteristic motivation to produce a certain desired end" is an 'action-guiding emotion' that is a "disposition to have a motive" (p. 137). It motivates us to act in such a way (i.e., develop and exercise virtue) that will be a "reliable success in bringing about that end" which is the target or goal of the motivation (p. 137).

For instance, the motivation for being generous in a situation that calls for generosity will involve the motive to actually act in a generous way by enriching someone else at one's own sacrificial expense, and that motive will be reliably successful in doing so when it is exercised. Regarding intellectual virtues, Zagzebski argues that their motivational basis is the "motivation for truth" which is the motivation to have "cognitive contact with reality" (p. 168). As such, having this motivation "leads us to be aware of the reliability of certain belief-forming processes and the unreliability of others" (p. 181). When a particular process (i.e., intellectual virtue) successfully contributes to achieving cognitive contact with reality upon being exercised and continues to do so reliably, the intellectually virtuous person recognizes this. The opposite

will also be true, however. Processes that are unreliable at attaining true beliefs, such as the reading and interpreting tarot cards and tea leaves, are duly noted by the intellectually virtuous person as dubious. As such, she disregards those ways of attaining beliefs because they do not do so reliably.[8] Such motivation, then, influences how we form beliefs.

In doing so, however, this motivational aspect of intellectual virtue also carries a significant social element that affects our epistemic behavior in terms of our epistemic responsibilities as cognitive agents within a community. Given our individual motivations for having true beliefs (i.e., 'cognitive contact with reality') this *ought* to motivate us to utilize those processes of belief-formation that are reliable at forming true ones, and hence eschew those that are unreliable. Our everyday practice of praising those (whether the local newspaper, our minister, a professor, or stockbroker) who get truth reliably and blaming those who don't manifests this. We blame (i.e., judge, look down upon, condescend) those who form beliefs that result from the exercise of unreliable processes (i.e., tarot cards, tea leaves, lucky guessing, etc.) since the use of these "reflect poorly" the agents that use them because it "shows a lack of motivation for knowledge" (p. 207). Furthermore, such blame doesn't only

[8] I should note here that even though Sosa and Zagzebski both understand their respective notion of intellectual virtue to have a 'reliability' component, it works differently for each thinker. On Sosa's account, reliability is the factor that gives a belief its justification: if a belief-forming process produces more true than false beliefs (i.e., if it's reliable), then its product belief is justified. For Zagzebski, although beliefs that are the result of intellectual virtue are reliable, such reliability isn't what confers justification on those beliefs, although it is a happy result of a belief formed out of intellectual virtue.

apply to those who exhibit a disconnect between a possible virtuous motivation and the procedures they exercise to bring it about, but it especially applies to those who don't have any virtuous motivation whatsoever and hence possess the vices of "dogmati[sm], closed-minded[ness], unfair[ness]" (p. 207). As such, this problem is not the individual's alone, but is a "social one since believers form a community, and his unreliability means that others in the community cannot trust his cooperation in their own pursuit of knowledge, nor can he trust himself in forming beliefs in the future" (p. 208).

This is an interesting parallel to Sosa's argument, noted in the previous chapter, regarding the need for mutual reliance between all believers. Sosa argued that we mutually rely on one another to figure out under what circumstances and conditions we are reliable, and under which of them we aren't reliable. Thus, we rely on one another to *exercise* the appropriate competences (i.e., Sosian-virtues) under the right conditions because we need this terribly, even for our own survival in some cases.

For Zagzebski, since knowledge is inherently a social phenomenon (and, is a necessary one for our survival and flourishing), we rely on one another for the appraisals of epistemic praise and blame we deserve in regard to our motivation for truth and the intellectual virtues we (ought to) develop and cultivate as the correct response to that motivation. Hence, we rely on each other for *possessing* intellectual virtues. So, even though there is a difference between Zagzebski and Sosa regarding what virtues are, this is a

place where they somewhat agree as to their importance at a social or macro level.

The question at this point, then, is how these features of Zagzebski's virtue theory figure into her definition of 'knowledge.'

Zagzebski, using the notion of the exemplar, understands a right or *moral* act as "what a virtuous person would or might do in a certain situation" (1996, p. 233). But she extends this to apply to *intellectual acts* as well: "the rightness of cognitive activity can be defined in a parallel fashion" as "what persons with the intellectual virtues would or might do in certain circumstances" (p. 233). 'Right cognitive acts,' then, are *beliefs* which are "cognitive states that […] result from the operation of the character traits of intellectual virtues" (p. 234). When a belief is the *result* of intellectual virtue, then it is justified: what *confers* justification on a belief is that it be formed as the result of intellectual virtue. Hence, 'justified belief' is defined in terms of intellectual virtue.

Knowledge, then, is "a state of belief arising out of acts of intellectual virtue" whereas the belief is also true (1996, p. 271). At this point, I should note that it may seem a bit odd for her account that 'justified beliefs' are the product of exercising virtue, and 'knowledge', then, is a 'justified belief' that is also true, and yet she doesn't define knowledge simply as 'justified true belief.' But Zagzebski has her reasons for not defining knowledge in this way, and I

try to clear this up in the next few pages. As such, there are some significant aspects of her definition that I need to point out.

First, it is important to note why, on this particular responsibilist account, knowledge is more valuable than mere true belief. For Zagzebski, even though knowledge is to be valued as a "possession [of the] truth," there's something even more valuable about the "relation between the knower and the truth" (p. 260). If knowledge arises *out of* the exercise of our intellectual virtues, then this suggests that it is something we work for successfully. It will arise out of the proper exercise of the intellectual virtues we possess (in our motivation to seek truth) which has actually achieved the truth that it set out to get in a way that avoids any possibility of getting it merely by luck. As such, Zagzebski avoids Gettier-anxiety by emphasizing that a belief's truth must be properly connected to being justified in a certain way, i.e., that the truth is reached "through an act of intellectual virtue" (p. 297).

Second, then, why is knowledge more valuable and desirable (and thus better) than having a merely justified belief? As Zagzebski points out, "justifiedness […] is a quality that even at its best only makes it likely that a belief is true" and from this "we call a person justified in having a belief because she has a property that (among other things) tends to lead her to true beliefs" (p. 268). As such, justifiedness itself doesn't guarantee truth. A belief's justification (i.e. being the result of an exercise of intellectual virtue) signals

that a belief has been formed in that admirable way, but it doesn't signal its truth necessarily or infallibly.

As such, knowledge is better than mere true belief or mere justified belief. For Zagzebski, it is a cognitive state above and beyond these primarily because of how much more restrictive and 'high-grade' an epistemic state it is. While a justified true belief's truth is attained through the exercise of intellectual virtue, a belief that is 'knowledge' is something more in that it manifests virtue in such a way that the agent is able to take a reflective and conscious perspective on it: it "requires the knower to have an intellectually virtuous motivation in the disposition to desire truth, and this disposition must give rise to the conscious and voluntary acts in the process leading up to the acquisition of true belief (or cognitive contact with reality), and the knower must successfully reach the truth through the operation of this motivation and those acts" (1996, p. 273). As such, the possessor of 'knowledge' is "cognitively integrated" in that she "has positive higher-order attitudes toward her own intellectual character and the quality of her epistemic states" (p. 275). She isn't merely an epistemic agent that forms true beliefs, but a *responsible knower* in the sense that she takes her epistemic role to be such that she could be blameworthy for it, and so conducts her epistemic life in such a way that exhibits the right kind of perspective on her own cognitive states. Such a high-order perspective on her own epistemic states means that "her belief structure is coherent, and she is aware of its coherence" (p. 275).

Given this particular analysis of knowledge, Zagzebski notes that reliabilists like Sosa who, in their attempt to define knowledge "to cover a multitude of states from the simplest cases of ordinary perceptual contact with the physical world […] to the most impressive cognitive achievements" (p. 262), define it "too cheaply" and to such an extent that it "does not really deserve the accolade knowledge" (p. 263). What sets 'knowledge' apart from "other good but lesser epistemic states" is that "the knowing state has a normative property that that state has in particular" as a result of intellectual virtue (p. 268). This particular 'normative property' is possession of "good intellectual character" (p. 275). The attainment of knowledge is not the mere following a certain *procedures*, *but that one is a certain kind of person*.[9]

At this point, some interesting parallels between her epistemological theory and Aristotle's moral theory come to the fore.

For one, virtue is primary. The ultimate state to be achieved ('knowledge' for Zagzebski's epistemological account, '*eudaimonia*' for Aristotle's moral account) is defined and made sense of only in terms of virtue after virtue has been discussed and defined. For Aristotle, once it is made manifest that virtue is that which makes people good, he thinks that then the notion of the highest good achieved by virtue – *eudaimonia* – makes sense.

[9] Of course, one question we should ask here is why knowledge must necessarily have this special character in order to rightly achieve this 'accolade' for Zagzebski. And, of course, the best answer may simply be that Zagzebski's goal is to present a virtue epistemology that is as Aristotelian-typed aretaic as it can be. Since virtue is paramount in Aristotelian moral theories, so must it also be on an epistemic theory that attempts to parallel it.

This is evident in his argument that people must be raised in good habits first in order to taste the 'good' and then develop their own motivation to pursue the moral virtues in order to make one's self good and thus achieve *eudaimonia*. Similarly, Zagzebski seems to follow this same explanatory structure, arguing that the intellectual virtues are worth having because they motivate their possessor to pursue having 'cognitive contact with reality.' So, when the possession of such virtue comes to fruition, that is, when they actually do the work in forming a belief that is true *because* they were exercised properly as the result of being 'cognitively integrated' into one's doxastic structure, that belief may attain the status of knowledge. As such, the highest moral and intellectual goods are defined in terms of virtue.

Second, Zagzebski's particular definition of 'knowledge' shares a particular affinity to Aristotle's notion of '*eudaimonia*' in that both may share somewhat arbitrary elements. Zagzebski admits that there is a point in her epistemological theory where "we find an exception" in her effort to structure it on Aristotle's moral theory because "no concept in ethics is exactly comparable to the concept of knowledge" (p. 272). Knowledge is that ultimately desirable intellectual state that one would want one's beliefs to achieve and this would seem to be a claim that *everyone* would agree to. But does that mean that everyone would agree to define knowledge in the same way? No. Zagzebski's preference, as she admits, is to describe its necessary conditions in such a way that doesn't allow it to be "gained too cheaply" (p.

263). But, again, this is her own understanding of knowledge as she sees its connection to virtue (and thus provides an argument to hopefully prove as much). Similarly, Aristotle notes that there will be no unanimous agreement as to what the content of *eudaimonia* actually is and how everyone will define it. To some it is mere pleasure. To others, including Aristotle, it is life in accord with virtue (1095a15-20). Yet, even though everyone does not agree as to what *eudaimonia* consists in, they do agree (given its broad definition of 'happiness' or 'the best life possible') that it is the most valuable and desired kind of life.

This is further manifested in how Zagzebski takes seriously Lorraine Code's claim that we are responsible for our intellectual actions and inactions such as we are for our moral ones, and then goes on to make the further claim that our intellectual responsibility is actually an aspect of our broader moral responsibility and thus of our character. As she claims, "a sense of responsibility for our own knowledge is as important for a sense of self-identity as a sense of responsibility for our own acts (p. 261). But even though she notes that "epistemic evaluation just is a form of moral evaluation," she also readily admits that she "cannot offer anything that would qualify as an account of the moral to support this claim," except to say that it probably isn't problematic since "there is no plausible apriori theory of the moral to start with" anyway (p. 256). Apparently, then, her theory of virtue, especially her definition of knowledge that relies on it, rests on something else for its

cogency and plausibility, possibly the influence of a macro-context (more on this in the next section).

One final parallel is that whether one wants to become a Zagzebskian knower or an Aristotelian *phronimos*, either status is a very arduous achievement. On the one hand, they both require an agent to go through a demanding process. Virtue is an achievement, after all, and not an accident. On the other hand, part of this achievement is that one's virtue is perfectly stable and secure. For Aristotle, being a virtuous person is such that virtue flows from a "firm and unchanging character" (1105a30) such that no matter the situation one encounters, she is "truly good and wise" and "always makes the best of circumstances" (1101a1-5).

Similarly, Zagzebski argues that a virtuous epistemic agent is *consistent*: "even when the motivational component of a virtue is generally related to success, we do not call a person virtuous who is not *reliably successful herself*" since just "merely being motivated to act is not sufficient" for being virtuous (p. 177). What distinguishes this person is that such virtues are "entrenched in a person's character" which will be necessary when "encounter[ing] resistance" in order to "reliably withstand the influence of contrary motivations when those motivations do not themselves arise from virtues" (p. 178).

4.4. Responsibilism and Contextualism

So, how is the responsibilist epistemology that I have outlined in the previous section contextualist in the way that I outlined it in chapter one? There, I claimed that regarding our epistemic goals there exists a *macro*-context that influences/determines what the standards are for practicing that goal-achievement in various *micro*-contexts or situations in which the goal is aimed at. The question I seek to answer, then, is whether or not responsibilist epistemologies – in which intellectual virtues are the primary evaluative property – rely on a macro-context in which standards are determined for attributing knowledge in various micro-contexts.

On the one hand, Lorraine Code's argument is thoroughly contextualist in both senses and I tried to point this out as I discussed her theory in §4.2. She relies on an understanding of the larger social context that has a role not only in making knowledge possible but also in that it attributes praise and blame to knowers given the extent to which they exercise intellectual virtue in the various situations they encounter. Code fits the mold of macro and micro, but her epistemology sets the groundwork for more details to be filled in, which is what Zagzebski does.

So, is Zagzebski's responsibilist epistemology macro *and* micro contextualist? If so, what makes it so? I argue that Zagzebski's *virtue* epistemology can indeed be interpreted in the contextualist way I have been discussing for many of the same reasons that I think Aristotle's basic ethical

framework can be reinterpreted in this light as well (as I discussed in the previous chapter).

To quickly review, let me summarize my discussion from the previous chapter. Aristotle's ethics might be conceived as having a *macro*-contextualist element because of how the *exemplar*, that standard which plays the all-important normative role for understanding and identifying virtue, is itself a sort of social (macro) construction. It may be conceived as having the micro-contextualist element, then, because the norms or standards that determine what is the proper particular manifestation of a virtue in one situation will differ from what that will be in other various situations. Furthermore, the macro influences the micro. How does an agent come to know which manifestation of generosity, for instance, is the correct one for a particular situation? By using one's 'capacity to reason,' which for Aristotle is the rational activity of trying to figure out what the virtuous exemplar would do in that particular micro context, i.e. how that exemplar would manifest virtue in that situation.

Similarly, Zagzebski's theory has an affinity to this sort of framework, but in an epistemological way. But on her account the exemplar isn't necessarily a construction of the privileged class or upper echelon of society from which the determination of norms emanates. Rather, she seems to conceive of the exemplar as a construction of the practitioners of the practice of epistemology. In this way, she seems to conceive of epistemic practice as a

sort of MacIntyrian practice where those epistemic agents of a particular macro-context are part of the group of people that enjoy the goods internal to that practice: being part of an epistemic community. Such a 'good,' then, is much broader than merely attaining the external goods of true beliefs (it is possible to 'cheat' and achieve these like we may cheat and achieve successful moves in chess). But truly being part of the epistemic community allows one to gain an understanding and perspective of the world from the vantage point of identifying with that community. And, so, this can *only* be achieved by engaging in epistemic practices such as basic responsible inquiry that gets cashed out in assessing and criticizing beliefs, for instance. In this way, truly being epistemically responsible and maintaining good intellectual character is a good internal to the virtuous practices of those in that community itself. Such practices, then, are those which the practitioner exercises and hence manifests in particular ways in the various micro contexts she encounters. As such, the exemplar is macro-determined but is practice-discovered, and this all revolves around the epistemic community itself. Let's consider this in a bit more depth.

First, consider the explicit role that the ***macro-context*** plays in Zagzebski's epistemological theory.

The most prominent *macro* part of her account regards how she notes that the intellectual virtues are to be learned: they are dependent upon having an *exemplar* to imitate. Consider her claim regarding this:

> I propose that the stages of learning the intellectual virtues are exactly parallel to the stages of learning the moral virtues as prescribed by Aristotle. *They begin with the imitation of virtuous persons*, require practice which develops certain habits of feeling and acting, and usually require an in-between state of intellectual self-control (overcoming intellectual *akrasia*) parallel to the stage of moral self-control in the acquisition of a moral virtue. In both cases, the *imitation is of a person* who has *phronësis*. (Zagzebski 1996, p. 150; *emphasis mine*)

Zagzebski argues that this is the case because there is no clear-cut rule-governed formula for developing and applying intellectual virtues. Interestingly, this reminds us of our frustration with Aristotle in the previous chapter regarding how one is expected to possess and manifest moral virtue as the intermediate between two contrary ways in which it can go wrong. What we ultimately got there was not a formula for finding this intermediary, but a *model* for how to go about figuring it out: do what the virtuous person would do, and to do this, you have to have a proper perception of the situation and then exercise *phronësis* in order to figure out what is the most appropriate way to manifest the virtue for that situation. As such, *phronësis* plays a similar important role here since

> no set of rules is sufficient to tell us when to place intellectual trust in the reliability of another, or what a person with intellectual courage, perseverance, or discretion would do, and so on. For this reason imitation of the person with *phronësis* is important for acquiring both intellectual and moral virtues. […] One learns how to *believe* the way she should rather than the way she wants in a way parallel to her learning how to *act* the way she should rather than the way she wants. […] She learns such intellectual virtues as open-mindedness, the ability to recognize reliable authority, and the ability to think up good explanations for a complex set of data by *imitating persons who have these qualities to an exemplary degree*" (Zagzebski 1996, p. 150-1).

Furthermore, that there is a desirable exemplar also suggests there is an anti-exemplar, the kind of person *not* to imitate because his epistemic activity expresses epistemic vice rather than virtue. Such a person may be Robert Gosse whom Code illustrated as someone who was intellectually blameworthy because of his lack of intellectual virtue regarding certain beliefs. This might also be others whose intellectual powers are motivated out of desires other than attaining truth such as "that some particular belief be true, […] to hold on to beliefs, […] that one's previously published views not be proven wrong" (p 146). Such desires, then, motivate the cultivation of defective traits such as "intellectual pride, negligence, idleness, cowardice, conformity, carelessness, rigidity, prejudice, wishful thinking, closed-mindedness, insensitivity to detail, obtuseness, and lack of thoroughness" (p. 152). In such cases, the exemplar of intellectual viciousness is a model to be avoided since it is the antithesis of what it is to be virtuously motivated to attain truth and thus the model of chronic intellectual blameworthiness.

But this forces us to ask a basic a question: why is it that a particular conception of an exemplar (in terms of its particular characteristics) is the one to be modeled and another conception of it (in terms of its particular characteristics) is the one to avoid? What privileges these particular conceptions in their particular ways such that makes one the accepted norm and the other the anti-norm? It seems that for Aristotle, what justifies his particular exemplar as being the accepted norm is, simply, because it *is* the

accepted norm (for better or worse) in his particular macro-context. It is what his readers already will have agreed to and understood as such. Is there something along these lines also in Zagzebski's analysis of how her particular exemplar – the one who exhibits truth-conducive intellectual virtues – has come to be the exemplar?

It's not as clear for Zagzebski if this is the case, that is, if the exemplar is such because it is merely accepted as such. Yet, for her, the exemplar seems to exhibit a macro-context's best ideas about what highly virtuous persons from our own experience – morally and intellectually – would look like and thus setting a standard for moral and intellectual behavior.

Interestingly, she suggests this, not in her 1996 work from which we get her epistemic theory but in a more recent (2011) article where she argues for the plausibility of the practice of positing moral exemplars. There, she argues that our identification of an exemplar is an instance of what Hilary Putnam and Saul Kripke were trying to get at in their way of accounting for our practice of 'direct reference' regarding our use of natural kind terms in language. Kripke and Putnam argued that a term like "gold," for instance, is useful as a way to correctly and intelligibly refer to the substance, like pointing to it and calling 'that' substance "gold," without having to know its deeper metallurgical structure or character. The term "gold," then, is "whatever is the same kind of thing or stuff as some indexically identified instance" (Zagzebski 2011, p. 50).

As Zagzebski notes, the Putnam/Kripke account was useful in that it helped to answer the question of how we can accurately describe things when "often we do not know the nature of the referent, and yet we know how to construct a definition that links up with its nature" (p. 50). As such, this proposal in the philosophy of language "began a revolution in semantics because it means that competent speakers of the language can use terms to successfully refer to the right things without going through a descriptive meaning" (p. 50).

While there may be some dis-analogy here (after all, gold is a substance that does indeed have an exact scientific definition which spells out all of the necessary and sufficient conditions that must be satisfied for a substance to be 'gold'), Zagzebski's point seems to be that this theory of reference gives some sense of an explanation as to how people can pre-theoretically (yet accurately!) refer to objects without knowing their scientific definition or description. Arguably, this may be partly because there is historical stability in calling gold 'gold,' such as if we think about this sort of description as part of a human narrative going back several hundred years, despite whatever variability might have happened over time regarding the exact description of it. But for Zagzebski, picking out the exemplar as such, for ordinary people, happens in a very similar way.

Zagzebski thinks that our understanding of moral concepts and actions are similarly rooted. Consider the following rather lengthy quote which pretty

much speaks for itself on how Zagzebski conceives of the exemplar as the standard-bearer that is macro-contextually constructed:

> Basic moral concepts are anchored in exemplars of moral goodness direct reference to which are foundational in the theory. Good persons are persons *like that*, just as gold is stuff *like that*. Picking out exemplars can fix the reference of the term 'good person' without the use of descriptive concepts. It is not necessary that ordinary people engaged in moral practice know the nature of good persons – what makes them good. In fact, **it is not necessary that anybody know what makes a good person in order to successfully refer to good persons**, any more than it was necessary that anybody knew what makes water water to successfully refer to water before the advent of molecular theory. We need not associate any descriptive meaning with 'good persons,' and users of our language can successfully refer to good persons even when they associate the wrong descriptions with the term 'good person.' As with natural kinds like gold and water, **people can succeed in referring to good persons as long as they, or at least some people in their community, can pick out exemplars**. Practices of **picking out such persons are already embedded in our moral practices**. **We learn through narratives** of both fictional and nonfictional persons that some people are admirable and worth imitating, and **the identification of these persons is one of the pretheoretical aspects of our moral practices** that theory must explain. Moral learning, like most other forms of learning, is principally done by imitation. Exemplars are those persons who are most imitable, and they are most imitable because they are most admirable. **We identify admirable persons by the emotion of admiration, and that emotion is itself subject to education through the example of the emotional reactions of other persons**. I am proposing, then, that the process of creating a highly abstract structure to simplify and **justify our moral practices** is rooted in one of the most important features of pretheoretical practices we want to explain, **the practice of identifying exemplars**, and in a kind of experience that most of us trust very much – the experience of admiration, shaped by narratives that are part of a common tradition.[10] (Zagzebski 2011, 51-2; **emphasis** mine)

The bottom line, then, is that in defining moral goodness or right behavior, we do so "via indexical reference to a paradigmatically good person" and this is something that we already collectively practice (maybe even in the sense of a MacIntyrean practice) (2011, p. 55). As such, identifying the

[10] Such historical examples of this she points to are Saint Francis of Assisi, Confucius, Jesus Christ, etc. (p. 53).

exemplar involves a different (less arduous) sort of inquiry (i.e., the ability to pick it out simply as part of our collective and historical practices) than one that would require us to come up with the necessary and sufficient conditions for what makes an exemplar. I think the point is not that the former sort of description is the only one that works (after all, surely we're better off in terms our epistemic value for having very thorough descriptions of things, especially of the exemplar!) but our being macro-contextualized in a particular community without having such an arduous definition of its exemplar seems to suggest that the former sort of description is the only one we need. After all, water is *necessarily* H20. But we don't need to know H20 in order to get along in the world as successfully hydrated beings. From identifying the exemplar as such, then, our virtues – that is, what we take to be virtues – are generalized into maxims for good moral behavior that is 'good' because it makes one more like the exemplar.

So, if this plays an explanatory role in the realm of moral virtue, and if Zagzebski is right that intellectual virtues are merely a subset of moral ones, then this also plays a similar explanatory role for intellectual virtue as well. As such, the standards and norms for good epistemic behavior which are the various virtues required for attaining knowledge really have their basis in what we, as a broad epistemic community, find to be admirable and desirable in others we identify as intellectually virtuous persons. And, so, they are those whom we identify or collectively 'point to' as examples of the intellectually

virtuous exemplar. As such, the discovery (and prescription) of the various virtues would seem to be a result of our analysis of what makes the exemplar the admirable model we wish to try to exemplify. Of course, why the particular intellectual or epistemic exemplar is admirable is because she is the model of someone who achieves knowledge (after all, she's the exemplar). But why does she have knowledge? In answering this question, it is said that the exemplar possesses virtues that are not merely conducive to forming successful beliefs but importantly instrumental (and even necessary) in attaining for those beliefs the positive epistemic status that we all aspire to possess. Such 'virtues' that are possessed by the exemplar that is identified and admired become the norms that govern our epistemic practices.

Also, Zagzebski's virtue responsibilism is ***micro-contextual*** in that once the exemplar has been identified (i.e. pointed out, constructed, accepted) at the macro level, then we have that which will determine the rules and norms that will have to be satisfied in various situations in order for an epistemic performance to count as virtuous in those situations.

On one hand, that this must be the case on Zagzebski's account is somewhat obvious from what has been presented so far. The role of the exemplar is to provide the norms for good and acceptable epistemic behavior for any particular situation because her possession of those norms is thought to be something we admire and thus desire to have for ourselves. Some situations may require the exercise of a particular epistemic virtue, say 'open-mindedness'

when encountering a new claim that challenges one that I have strongly held for a long time. But in another situation, such as one in which a student challenges my (i.e., her philosophical superior) ideas in class, it may call for me to exercise more 'intellectual humility' than 'open-mindedness,' or maybe I will still be responsible for behaving in an open-minded way, but not in exactly the same open-minded way as in the other situation. After all, I should always maintain and exercise some degree of open-mindedness – more so when reading the latest research on preventative measures for cancer, less so when my friend offers today's prophecy via her morning's coffee grounds. But without an exemplar to point to as the model for proper intellectual behavior to which I am striving to conform my intellectual life (to be virtuous, after all), I wouldn't be able to make this assessment in an objective and normative way. As such, the role of the exemplar necessarily implies the existence of micro-contexts to which the norms and standards of the exemplar are to be applied.

But let's also consider what Zagzebski explicitly says about the micro-context.

First, just as the exercise of virtue requires experience on Aristotle's account, so does it for Zagzebski as well. One who acquires and exercises virtue doesn't do so at the "flip of a switch," but due to her "series of experiences of interaction with the world around her" (Zagzebski 1996, 117-20). As such, the refinement of our virtue is achieved through exercising it in

particular situations. In this sense, then, the practice of virtue in micro-contexts plays an educational role that makes us more virtuous agents.

Secondly, consider the aspect of Aristotle's theory that I noted in the previous chapter regarding how the virtuous person possesses *phronësis* which then equips her to know exactly which virtue would be called for in a situation and to what extent it should be manifested. On Zagzebski's account, *phronësis* has the same role. In particular, she notes that it is the "higher order virtue that governs the entire range of moral and intellectual virtues" (p. 229). As such, it is "necessary to make sense of […] justified belief in a virtue theory" because it is a necessary component of determining "the [intermediate] *at the time of action*" or 'in a particular situation' (p. 220; *emphasis* mine). Even though "a virtue is always a good thing for a person to have […] there is a complication in that equal degrees of a virtuous trait are not always associated with equal degrees of internal good in the agent" and thus "there is no determinate degree and kind of activity exhibited by a virtuous person that is uniformly good-making" (1996, p. 96-7). For Zagzebski, then, an intellectual virtue's exercise will necessarily look (i.e. be manifested) differently given the particular micro-context in which it will be exercised as opposed to another micro-context.

With this in mind, Zagzebski notes some important reasons why *phronësis* is necessary. Primarily, these have to do with its "mediating function" in the sense that it allows a virtuous person to "be aware of the facts that

contextualize the belief in question" which is to say that it allows her to operate in such a way that picks out all the relevant facts of a situation, about herself, and hence what the proper manifestation of an intellectual virtue is, which will then result in a justified belief as a result of the operation of such virtue (pp. 245-6).

Phronësis allows the agent to act virtuously in spite of variability and personality differences between different agents. So long as they each possess *phronësis*, they can take their own personalities into account in relation to any particular micro-context they find themselves in and come up with the correct, virtuous intellectual action.

> Some persons are intellectually successful by careful plodding, whereas others do their best work when they indulge in exuberant intellectual impetuousness. These individual differences seem to strain the distinction between virtue and vice since the mean *relative to a particular individual* may lean very far in the direction of one vice or its contrary. Still, it is fair to say that human thinkers should be neither too impetuous nor too slow and plodding, neither too steadfast in their views nor too intellectually independent, etc. There is such a thing as too thorough, too careful, too attentive to detail.[11] But *what counts as too much or too little may very from person to person, as well as from one intellectual context or discipline to another* (Zagzebski 1996, pp. 96-7; *emphasis* mine).

Zagzebski further develops this line:

> For example, in the case of intellectual carefulness, it takes practical wisdom, not speculative wisdom or some other virtue, to tell how much evidence is enough to support a belief; clearly, it will differ from case to case. The same point goes for perseverance, courage, and autonomy. We need to make choices in intellectual inquiry, just as much as in deliberation leading to

[11] One example Zagzebski cites from Lawrence Serne's *The Life and Opinions of Tristram Shandy* in which the character 'Uncle Toby' is described as possessing the intellectual vices (i.e. defects) of excess: "excessive attentiveness, thoroughness, diligence, perseverance" which are traits that "have sometimes been called, with irony, 'virtues in excess' (Zagzebski 1996, p. 196).

> action, and the extent to which we should persevere in a line of inquiry and answer attacks from others is a matter of judgment. We need the same kind of judgment to **strike the proper balance** between relying on the authority of others and relying on our own intellectual powers. […] It takes *phronësis* to know how persevering one should be to be persevering, how careful one should be to be careful, how self-sufficient one should be to be autonomous, and so one. The first theoretical need served by *phronësis*, then, is to determine the mean between extremes in those cases in which the virtue is a mean. (Zagzebski 1996, pp. 220-1; **emphasis** mine)

And further:

> A single set of circumstances can easily be an instance of more than one such feeling or activity, and in such cases more than one virtue applies. The features of a situation that are relevant for one virtue may not be for another. […] If we are to avoid an excessive and unworkable fragmentation of value, there needs to be some virtue that permits a person to sift through all the salient features of the situation – that is, all those features that are pertinent to any of the virtues – and to make a judgment that is […] the judgment of the virtuous person. […] Mediation between intellectual courage and intellectual humility or fairness is necessary just as much as mediation between moral courage and humility or fairness. A person needs to know at what point to make an intellectual commitment just as much as she needs to know when to make a moral or a personal commitment. But intellectual commitment can run up against the virtues of intellectual caution, thoroughness, or fairness to the views of others. Knowing what to do in these cases is not simply a matter of having a combination of virtues in question. **A virtue is needed that permits a person to see the big picture, and this will be a virtue that balances such virtues as carefulness, thoroughness, and fairness with perseverance and commitment**. Intellectual virtues, then, need a mediating virtue […]. As far as I can tell, the only candidate for such a virtue in the philosophical tradition is *phronësis*. (Zagzebski 1996, pp. 221-3, **emphasis** mine)

Why *phronësis* is needed, then, is because it is that which makes it possible to determine what to do in a given micro-context. In Zagzebski's theory, it would seem that one would suffer from utter frustration in how to form a belief based on the model of the exemplar if that's all we had, just an exemplar that [prescribes norms and rules, and not a way to make that exemplar work for us in terms of some kind of framework for applying it to

our own particular situations. As such, to claim that 'we just need to be intellectually virtuous' in order to have justified beliefs is a bit underwhelming. Indeed, Zagzebski argues that this frustration is characteristic of philosophy itself.

> The problem here is that virtuous procedures underdetermine action […] Even when we list all of the known procedures used by a virtuous person, there is still an indeterminacy with respect to most beliefs. […] This is to say, not that most persons are irrational, but rather that rationality is underdetermined by procedures. […] On a typical day we encounter propositions about such matters as the consequences of various fiscal policies, the guilt or innocence of an accused murderer, the trustworthiness of particular politicians, the greenhouse effect, the weather forecast. To be honest we must admit there is no adequate evidence for most positions on these matters. […] ***Phronësis*** **is needed, then, [because] many human activities, whether of the overt kind traditionally handled by ethics or the internal activities of thinking and forming beliefs, can be neither fully described nor evaluated in terms of the following of a set of known procedures or rules. Good judgment is required in all areas of human activity, including the cognitive**. Persons with practical wisdom learn how and when to trust certain feelings, and they develop habits of attitude and feeling that enable them to reliably make good judgments without being aware of following a procedure. There is a very strong element of inclination in most beliefs, even in the beliefs of those persons most intellectually practiced and aware. The difficult part is to train the inclinations themselves to reliably produce the desired end – in the case of intellectual activity, knowledge. (Zagzebski 1996, pp. 224-6; **emphasis** mine)

Ultimately, what Zagzebski is implying is that *phronësis* really, then, is a way to describe how we attempt to make our behavior line up with that of the socially constructed exemplar. *It is only in our best judgment* that we make ourselves mirror the exemplar since the notion of the exemplar itself is quite broad and vague at best. So, Aristotle's notion of the 'capacity to reason' seems to equally apply here: we need to be able to judge how to best exemplify the

exemplar in our intellectual behavior. And *phronësis* is that capacity that allows us to do this.

Finally, then, we see in Zagzebski's account a possible affinity to the kind of contextualism offered by DeRose and noted in the first chapter. When it comes to accounting for the epistemic norms and standards that are accepted and practiced in a community, at the macro-level such norms are identified, and they are applied – although variably and differently – in the various micro-contexts. Interestingly enough, one might be inclined to resist the notion that different situations require a different application or exercise of virtue. After all, we tend to think of virtue as something quite firm and concrete (I will discuss this particular aspect of virtue more in the next chapter). But the Aristotelian conception of virtue affirms that its application is not uniform among the various situations in which it should be applied. As noted in the previous chapter and in this one as well, intellectual virtues – the epistemic standards and norms for those micro-contexts in which we form beliefs – that are exercised in various situations are manifested differently, that is, they look different in different situations. Interestingly, this shares a striking affinity to attributor contextualism's explanation as to why some micro-contexts (such as low-stakes ones) call for a particular set of epistemic standards while others (such as high-stakes ones) call for a different set.

Zagzebski's virtue responsibilism offers us an account where such standards – intellectual virtues – are understood as being applied in different

degrees (particular intermediates/manifestations) as well differences in kind (e.g., open-mindedness or intellectual humility) depending on the micro-context itself and what it particularly calls for. And she then gives us a macro-contextual reason as to how this is accomplished: by conforming one's intellectual behavior to the norms and standards of the macro-contextually constructed exemplar, i.e. exercising the intellectual virtues.

Chapter Five

The Generality Problem and the Situationist Critique

> *In the same way that the situationist challenge depends on whether moral virtues are coarsely individuated (e.g., honesty, courage, modesty) or finely individuated (e.g., honesty while watched by fellow parishioners, courage in the face of rifle fire, modesty before peers while in a good mood), the generality problem depends on whether intellectual virtues are coarsely individuated (e.g., curiosity, creativity, intellectual courage) or finely individuated (e.g., curiosity while in a good mood, creativity after being given candy, intellectual courage in the face of non-unanimous dissent.*
>
> - Mark Alfano (2012, p. 233)

5.1. The Problems with Virtuous Processes and Virtuous Character

Thus far, I have argued that two epistemic virtue theories – virtue reliabilism as articulated by Ernest Sosa and virtue responsibilism as articulated by Linda Zagzebski – are contextual theories. What I mean by calling them 'contextual theories' is that they inherently contain the two dimensions I identified in chapter one that makes up the core of attributor contextualism. They both make use of *attributor* factors (which I have described as the macro-context where the community authoritatively sets the norms/rules for particular epistemic situations). They also make use of *subject* factors (what I have called a micro-context where the norms/rules that are determined by the

macro-context are played out in the particular situations themselves that vary among themselves in terms the degree to which the norms/rules are required). For virtue reliabilism, at least in its Sosian form, the macro-context is that context which determines the level of reliability that is sufficient for certain kinds of epistemic performances' justification relative to particular circumstances or micro-contexts. For virtue responsibilism, at least of the Zagzebskian variety, the macro-context determines the conception of the exemplar who serves as the model or norm of what 'good' (i.e., cognitive truth-seeking) epistemic behavior is, and micro-contexts are those situations in which that norm is applied. This is why some situations might call for a certain virtue to be manifested in a certain way while in other situations in a different way.

As such, the role of the *macro* in determining the standards for the *micro* seems to arise out of the needs we have as epistemic beings. As a community, we need our epistemic practitioners to be successful ones. And so, Sosa identified our individual 'reliability' (grounded in the reliability of our competences or 'virtues') as that which the community is interested in us possessing and exercising. For Zagzebski, being epistemically 'responsible' is what the community is interested in, and a complete model of such responsibility takes form in the admirable traits or 'virtues' that come together to form the intellectual exemplar.

Each epistemic theory, however, faces significant problems. Some have even argued that their problems are quite insurmountable and have no good prospect for any sort of resolution, much less any solution.[1] The problems for virtue reliabilism and responsibilism are the *generality problem* and the *situationist critique*, respectively. The former problem points out the inherent complication that results from making reliability *the* justification-conferring factor because since a belief's justification would then depend on the reliability of the process that produced it, we would need to be able to *describe* that process as accurate as we can (not too broad, not too narrow) in order to gauge just how justified that belief actually is. The latter problem is directed at the fact that since virtue responsibilism requires epistemic virtues to be character traits that are 'entrenched' in a person's character so that the intellectually virtuous person cannot not exhibit those virtues, it faces a serious challenge from social psychological experiments that show that the reason why people get stuff 'right' is not due to their virtue, but due to seemingly irrelevant features of situations themselves.

[1] I might agree with this charge, if only because it seems that the trend is that philosophical problems are ultimately answered and solved only when a significant amount of the philosophical community attach themselves to a particular response and declare it 'solved!' If this is what it takes, then the problems this chapter deals with are far from achieving that status.

5.2. The Generality Problem

Alvin Goldman was the first to come right out and tell us that his state-of-the-art reliabilism was flawed, although he didn't make a huge issue of it at the time. Remember, on his theory, "the justificational status of a belief is a function of the reliability of the process or processes that cause it where […] reliability consists in the tendency of a process to produce beliefs that are true rather than false" (1992, p. 113). But as he noted, "a critical problem" of this will be determining what is "the degree of generality of the process-type in question" since however we describe or identify a particular process, whether more broadly or narrowly, we will inevitably pick a particular level of generality within that spectrum that "will determine the degree of reliability" (p. 115).

The reason why this is such a 'critical problem' is because Goldman's offering of 'reliability' as a new way to think about justification meant that the latter was dependent upon the former – the *more reliable* a belief forming process is, the *more justified* the output beliefs produced by that process. So, if we can't get an accurate descriptive handle on what the actual process is that forms beliefs in order to determine its reliability, then how can we actually gauge those beliefs' justification? Richard Feldman thinks we can't do this and probably never will since "it is only by describing processes in an arbitrary and ad hoc way, on a case by case basis, that the reliability theory even seems plausible" (1985, p. 159). For Feldman, a principled way of determining reliability is what reliabilism needs but cannot have.

To understand this problem more in depth, we need to consider aspects of what it means for there to be a reliable belief-forming process. Here we must distinguish between *types* of processes and *tokens* of a process.

A process *type* is any particular species of belief-forming process, such as a visual-perception belief-forming process (that forms beliefs about the things I see) or a inferential-logical belief-forming process (that forms beliefs about the propositions I logically infer from other beliefs/propositions), or a memory-operational belief forming process (that forms beliefs about or from the memories I have), and so on. *Tokens* of a process are the particular instantiations of that process at any one time, understood as the "dated sequence of events that results in a belief" (Feldman 1985, p. 159). A "token event sequence," then, is merely a particular occurrence of that process's operation that produces an output belief (Conee & Feldman 1998, p. 2).

It might be helpful to think of this as analogous to an automobile assembly line in my hometown of Detroit. There are different *types* of assembly lines that produce automobiles. One produces the *Jeep Cherokee* (Toledo Assembly), another produces the *Chrysler 200* (Sterling Heights Assembly), another produces the *Ram Truck* (Warren Truck Assembly), another produces the various engines that will be installed in all of these vehicles (Mack Engine Plant), while another will assemble the transmissions that will be linked up to those engines (Fort Wayne Transmission Plant). A *token* of any of these various processes, then, will be any particular instance of

the operation of one of those processes, such as the particular time when a particular *Chrysler 200* is built from start to finish. As such, that *200* is a *product* of a particular instantiation of the *Chrysler 200* automobile-forming assembly line process at Sterling Heights, Michigan.

What is it, then, that we must properly look to and analyze in order to determine reliability (or unreliability)? Surely, not the tokens themselves. It makes no sense to say that tokens are reliable since a token is merely a particular instance of a process having operated at a particular time. But tokens are important for gauging reliability in that they inform us of the success (or failure) of the processes that produce them. The product (whether a car or a belief) of a token is either a success (for cars, they run – for beliefs, they're true) or a failure (lemons and errors). So, if upon evaluating the *totality* of token instances of a particular process type we find more successes than failures, then we must say that that the process *type* is 'reliable' since, after all, it has produced more successful token instances than not. Thus, the proper object of reliability is the process itself, the assembly line.[2]

[2] However, I do realize that in common parlance we more often refer to particular cars themselves, the *products* of process types, as either 'reliable' or 'unreliable.' As such, the analogy is an imperfect one. However, the reason why is because cars have a feature that beliefs don't: they themselves are a sort of process type in that the particular instantiations of their being exercised – what we would call a 'trip' however short or long – either gets us from point A to point B or not. When a car does this more often than not, it is reliable. However, even this is contextually determined in the sense that although a car that operates successfully 51% of time is technically reliable, it wouldn't be practically reliable, and we wouldn't attribute it as being reliable. Actually, it may be the case that we need a higher threshold of successful operation, maybe even 98% or 99%, for a particular car before we'll say that it's 'reliable.' More on this in §5.6.

Applied back to reliabilist epistemology, then, when a process type is a reliable one, this means that all of its products – the beliefs it forms – will inherit their justification in virtue of their maker's reliability. As Feldman notes, "it is the reliability of the process type responsible for a belief that determines its justification" (1985, p. 159). As such, the justificatory status of a belief should be easily discovered so long as we can measure how reliable a belief-producing process is. Just as we can account for the reliability of automobile assembly lines, then, we should be able to account for the reliability of belief-forming processes.

The trouble, however, is that in the realm of beliefs, as opposed to cars, this is not so easy. Two issues immediately emerge and highlight this difficulty.

First, if we're going to determine the justification of a particular belief in virtue of the reliability of the process type that produced it, can we actually pinpoint one particular process? Isn't it the case that "the specific token that leads to any belief will always be an instance of many process types"? (Feldman 1985, p. 159). After all, any particular *Chrysler 200* has been the result of many different processes: the Bosch plant that produced its plastic interior components, the plant in Fort Wayne that produced its transmission, the Mack Engine Plant that made the 4-cylinder motor that will power it, and the Sterling Heights Assembly Plant that takes all of those individual components and assembles them together to make the finished product you see rolling

down the highway. So which process type is the one we should gauge for reliability? After all, some beliefs are the product of a combination of all sorts of process types working independently as well as together, just like with automobiles, such as visual-perception belief-forming processes, inferential-logical belief-forming processes, memory-operational belief forming processes, etc. So, from which process type will such a belief gain its justification since it would seem that the process types will differ from one another in terms of their respective reliability? For the purposes of explanation, Feldman claims that it would seem that "one of these types must be the one whose reliability is relevant to the assessment of my belief" and so let's say that there is "some relevant type" whose reliability is the most salient one since it explains the belief's justification (1985, p. 160). But this leaves us with the task of figuring out what the relevant process type actually is.

Second, then, we're faced with the problem of *describing* what the relevant process type is for a particular belief produced by that type. This is a problem because how we describe a particular process will affect our evaluation of how reliable it is. And so, since a belief's justification is dependent upon the reliability of the process type that produces it, and since our description of the process will determine its reliability, how we describe that process will directly affect our evaluation of whether or not (or, at least, to what degree) its product-belief is justified. As Feldman notes, in trying to do this we can go wrong in either of two ways:

> In coming up with an account of relevant types, defenders of the reliability theory must be guided by the following point. If relevant types are characterized very narrowly then the relevant type for some or all process tokens will have only one instance (namely, the token itself). If that token leads to a true belief, then its relevant type is completely reliable, and according to [reliability theory], the belief it produces is justified. If that token leads to a false belief, then its relevant type is completely unreliable, and, according to [reliability theory], the belief it produces is unjustified. This is plainly unacceptable, and in the extreme case, where every relevant type has only one instance, [reliability theory] has the absurd consequence that all true beliefs are justified and all false beliefs are unjustified. We can say that characterizing relevant types too narrowly leads to 'The Single Case Problem.' A very broad account of relevant types of belief-forming processes leads to what we may call 'The No Distinction Problem.' This arises when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type. For example, if the relevant type for every case of inferring were the type 'inferring,' then [reliability theory] would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type (1985, pp. 160-1).

This problem, then, is the challenge to "provide an account of relevant types that is broad enough" and yet "not so broad" in order to avoid 'The Single Case Problem' and 'The No-Distinction Problem,' respectively (p. 161). As such, the way we describe a process must be accurate and exhaustive in order to truly get us a precise measure of its reliability.

Another dimension of this problem is that a *process* type is not merely the mechanism itself that produces product-beliefs but also includes the conditions or circumstances under which the mechanism operates. As such, the description of the process must also include not only a description of the mechanism itself (such as the assembly line in the car example) but also the conditions under which that mechanism works because describing those

conditions are going to help us figure out how reliable the mechanism is under some kinds as opposed to under other kinds.

As I noted in chapter two, Sosa thinks it is very significant that as an epistemic community we determine what are the 'appropriate conditions' for the exercise of our epistemic competences (i.e., virtues). This is so, partly, because we need to know how reliable (i.e., in terms of the acceptable degree or threshold) we are and hence what we should require and expect of each other under conditions that are normal as well as under conditions that aren't. Furthermore, such conditions themselves play an important role for Sosa in that a justified belief (i.e., a belief that is the product of a reliable epistemic virtue) that is true and has been formed by that virtue in normal conditions (conditions that are appropriate for its exercise) will be a case of knowledge so long as what makes that belief true is the proper working of the epistemic virtue bringing about its existence. And that virtue will seemingly not be able to work properly under conditions that it's not suited to operate under. Just as the local automobile assembly-line mechanism will not work properly (i.e., will not successfully produce excellently built automobiles) if its line-workers are all drunk or if the temperature inside the plant is so hot that the machinery overheats and can't keep up with other machinery that's working properly, so it is that my visual-perception belief-forming mechanism, for instance, will not work properly (i.e., will not successfully produce true beliefs) if I'm quite

inebriated or if my body temperature is so high that I become delirious, or if the air quality is too smoky within the scope of my visual field.

Robert Brandom notes that one of the insights of the reliability theory of justification is its emphasis on the importance of the conditions for epistemic performance, that "the difference that makes the epistemological difference" in cases where we're trying to understand a belief's degree of justification is "in the circumstances in which the belief was formed" (2000, p. 115). As such, even if we're inclined to think that a belief-forming mechanism itself is either reliable or unreliable, such evaluations will be inaccurate (and thus meaningless) without taking into account the kinds of conditions that the mechanism can properly work under as well as those in which it cannot. However, this adds another level of complexity to the problem of coming up with the proper description of the relevant process at work – in describing the *process*, we have to describe the *mechanism* and the *conditions* under which it operates reliably.

A very insightful way to illustrate this is the way that Brandom describes conditions as 'reference classes' for belief-forming mechanisms. Consider Goldman's famous barn façade example. Barney is enjoying a beautiful sunny day by joyriding on the country roads of Barn Façade County and admiring all of the barn façades (which he doesn't realize are mere façades) that he passes as he takes his leisurely drive. Coincidentally, he stops at the only real barn in the county (that is, the only one that is not a façade but

an actual barn) and forms a belief that he is viewing a barn. He has no idea that it is the only real barn in that county, nor does he have any awareness that all the rest are mere façades.

The question that arises about this case is whether Barney's belief that he is viewing a barn is a case of knowledge or not. On one hand, Barney "has a true belief, has good reasons for that belief, and stands in the right causal relations to the object of his belief. Surely one wants to say […] what he has in such a case is perceptual knowledge" (Brandom 2000, p. 114). But as Brandom correctly notes, "things are less clear as we describe the case further," that is, as we take note of the circumstances in which Barney formed his belief since he would have formed that same belief had he been viewing any of the barn façades throughout the county (p. 114). It is only by luck, a coincidence, that he happened to stop at the only real one and form a belief about *it*. As such, Barney's visual-belief-forming-mechanism was employed in abnormal circumstances, and as such, that particular process is an unreliable one and Barney's resulting belief is unjustified.[3]

[3] Of course, this is one instance where Sosa's distinction between Animal and Reflective Knowledge can come to the rescue. What Sosa argues is that the abnormality of the circumstances only affects reflective justification. As such, Barney is not reflectively justified – the conditions are such that he cannot be aware of the factors that make Barn Façade County full of barn facades. That being said, however, Barney is animally justified – so long as his visual-perception belief-forming mechanism is operating under conditions that are appropriate for its operation, and the process is reliable under those conditions, then Barney's perceptual belief about the barn is indeed justified. The luck involved, then, doesn't affect the perceptual belief. It's a correct and justified belief no matter what. The presence of the luck, however, precludes Barney from ever being able to take a reflective perspective on his situation there.

What's interesting about this case is that the *circumstances* behind the belief – that Barney is in a county where all barns are mere facades except for one barn – direly affects any possible justification for any such beliefs in that county since it means that Barney's visual-belief-forming-process-about-barns is an unreliable one there. But what's even more interesting is what happens when the circumstances themselves are descriptively changed, that is, when the 'reference class' is broadened. As Brandom suggests, what if Barn Façade County is located in Real Barn Province where all barns (except for those in Barn Façade County) are real? Then, the conditions are such that Barney is a reliable perceiver of barns. But, what if that Province is located in the United Republic of Fake Barns? And, what if that nation is located on Genuine Barn Continent, Incognito Barn Earth, or Authentic Barn Galaxy? (Brandom, p. 115-6).

What this suggests is that all of these descriptions could be true at the time that Barney formed his belief about that one measly barn in Barn Façade County. But here's the kicker: "which is the correct reference class?" (p. 116). Which one is the proper one to note as part of the description of Barney's belief forming process here? Since "the reliability of the belief-forming mechanism […] varies depending on how we describe the mechanism and the believer," we can get either result – that he is reliable or unreliable – merely by choosing one reference class or another to use in our description of the process. But which one is the most relevant one here? How would we privilege one

description over another? What justifies our use of one reference class rather than any of the others?

Robert Hudson offers a similar example:

> For instance, a seemingly uncontroversial process for generating true beliefs involves a normal human being visually inspecting a mid-sized object in broad daylight and then forming beliefs on that basis. But now add the further fact that this human has just ingested a large quantity of hallucinogens – the process no longer looks reliable. Now further add the fact that the human has also ingested an effective anti-hallucinogen antidote. Again, the process seems reliable. (Hudson 2004, p. 193).

What this points out is that the "process of re-conceiving the reliability of a process appears to be interminable" since "depending on the detail with which one describes a process, it may turn out to be reliable or not" (Hudson 2004, p. 193).[4] So, which description is the best, most relevant one? In other words, at what 'level of generality' should we land on and why should we land on the one that we do? Arguably, nothing about the facts of this or Barney's case allows us to, in a principled way, pick one level or description at the exclusion of the rest. Picking a description without the use of any norm or principle would be to do it *ad hoc*, which is what Feldman argues ultimately dooms reliabilism as a theory of epistemic justification.

5.3. The Situationist Critique

As we saw in the previous two chapters, a mantra of virtue ethics and virtue responsibilist epistemology is that virtuous persons are those who are

[4] This point was also made by John Pollock in (1986).

motivated to exercise virtue and hence are motivated to possess them in order to accomplish praiseworthy actions (or beliefs). As such, virtues are the relevant descriptive or explanatory factor for why a virtuous person *acts* virtuously (when she does) as well as why we would appraise her as *being* virtuous: she possesses those traits which are entrenched in her character and have become part and parcel of who she is. Of course, we can point to proof texts for how this is the case, such as Aristotle's claim that the person who is truly virtuous "always makes the best of circumstances" (1101a1-5) and presumably does so because such action "proceed[s] from a firm and unchangeable character" (1105a30-35). And there is also Zagzebski's claim that an intellectual virtue is a "deep and enduring acquired excellence of a person" (1996, p. 137) that is an "entrenched quality" (p. 125) of one's character which is a "necessary feature of virtues because they are often needed the most when they encounter resistance" (p. 178).

The *situationist critique*, however, pokes holes at the notion that virtues are the essential explanatory factor of 'good' moral and intellectual behavior. Relying on the results from several social-psychological experiments, the critique's proponents argue that virtue is irrelevant for explaining and predicting moral behavior. At best, even if there is such a thing as virtue, it is not the case that good moral and epistemic behavior is explained in terms of virtue or *because* of virtue. At worst, there is no such thing as virtue.

John Doris was one of the first to bring to our attention this problem, ultimately arguing that virtue ethical theory is "descriptively inadequate" in that it doesn't capture what's really going on when people act 'morally' (Doris 1998, p. 505). Doris challenges both the folk understandings of morality that attribute moral behavior to virtue as well as the more sophisticated Aristotelian-styled virtue ethics that touts that "the person of good character is not easily swayed by circumstance" (p. 505). Within both kinds of perspectives, it is thought that people who possess such virtue are "dependable, steadfast, unwavering, unflinching," and so on (p. 505). Because of this, "character is expected to have regular behavioral manifestations" in that the virtuous person, even when faced with particular temptations or situational circumstances that might lead her to act contrary to character. So long as she possesses those deep-seated traits, she won't succumb to situational seduction since her 'virtue' is much stronger than that (p. 505). Virtue-ethics of an Aristotelian stripe, then, considers virtues as "robust" and "substantially resistant to contrary situational pressures in their behavioral manifestations" (p. 506).

But Doris's question is this: even if this is so, shouldn't we require empirical evidence for it? Doris thinks so. And he argues that the kind of empirical evidence we should require is an "observed behavioral reliability" which means that we should be able to see "behavior consistent with a trait or grouping of related traits across a range of relevant eliciting situations that may

vary widely in their particulars" (Doris 1998, p. 506-7). Virtue, then, to be accurately attributed to someone, should be the kind of thing that is consistently exhibited in her actions. No matter what kinds or quantities of situations she encounters, her virtue ought to be exercised and observed in each of them. Another way of putting this is in psychological terms, that the agent should exhibit "cross-situational consistency" in her behavior (p. 507).

But Doris argues that when we actually do the kinds of experiments that measure peoples' behavior and determine the reasons why they act as they do, we find that they repeatedly fail at exhibiting the kind of consistency in virtue that we uncritically attribute them as possessing and exercising. As Doris notes, "behavior is very often surprisingly unreliable" (p. 507). One example of this is from Isen and Levin's (1972) experiment that attempted to gauge what it was that actually motivated people to act in helpful ways toward others.

> Imagine a person making a call in a suburban shopping plaza. As the caller leaves the phone booth, along comes Alice, who drops a folder full of papers that scatter in the caller's path. Will the caller stop and help before the only copy of Alice's magnum opus is trampled by the bargain-hungry throngs? Perhaps it depends on the person: Jeff, an entrepreneur incessantly scheming about fattening his real estate holdings, probably won't, while Nina, a political activist who takes in stray cats, probably will. Nina is the compassionate type; Jeff isn't. In these circumstances we expect their true colors to show. But this may be a mistake […] The paper dropper [is] an experimental confederate. For one group of callers, a dime was planted in the coin's return slot; for the other, the slot was empty. Here are the results:
>
> | | Helped | Did Not Help | |
> |---|---|---|---|
> | Dime | 14 | 2 | |
> | No Dime | 1 | 24 | (Doris 1998, p. 504) |

This experiment and many others like it[5] seem to prove, for Doris, that *global* traits – those Aristotelian-typed, long-term, broad, enduring, and entrenched traits that are thought to reliably motivate 'virtuous' action in their possessors and thus justify (and are reflected in) *our* attributions of such traits to them – are ineffective at best, and non-existent at worst. The bottom line, then, is that they are irrelevant. They don't really do the kind of moral-explanatory work we often think they do. The more relevant explanatory factors for why people act morally good (or even bad) are the *local* features of their particular experience, that is, features of the *situation* itself that they must choose how to act in and how those features, then, come to bear on their actions.

In the example cited above, the relevant feature that affected behavior outputs was money – it had to do with whether or not the test subjects found a dime in the coin return slot. Whatever virtue or character trait they might have 'possessed' was unimportant. As such, Doris argues that 'situationism' is a "more empirically adequate conception of moral personality" (1998, p. 507). In short, situationism's three theses are

[5] Doris also cites the Milgram (1974) experiment and how that if character traits, such as compassion, were truly 'entrenched' and 'stable' and 'unfailing,' then it shouldn't have been the case that the subjects who were the 'punishers' in the experiment carried out the instructions to shock the 'learners' with fatal levels of electricity (all staged, of course, except the 'punishers' didn't know this at the time). Doris's response to this is that the punishers' "moral sensitivities appeared intact, but dispositions to act on them were overwhelmed by the demands of the experimental situation" such as "feelings of obligation, or perhaps, intimidation, generated by the experimenter's insistence on their continued participation in the 'learning experiment' (Doris 1998, p. 510).

> (i) Behavioral variation across a population owes more to situational differences than dispositional differences among persons. […] (ii) Empirical evidence problematizes the attribution of robust traits. Whatever behavioral reliability we do observe may be readily short-circuited by situational variation: in a run of trait-relevant situations with diverse features, an individual to whom we have attributed a given trait will often behave inconsistently with regard to the behavior expected on attribution of that trait. […] (iii) Personality structure is not typically evaluatively consistent. (Doris 1998, p. 507).

Since virtue is not morally meaningful, then virtue-attribution or trait-attribution is neither a precise nor adequate way to predict or appraise behavior. A prediction that someone will act a certain way in a particular situation because they have been attributed with possessing a particular virtue will be an inaccurate prediction. A more meaningful kind of prediction, then, will be one that takes situational features into account and how they might come to bear on the kind of behavior they motivate and instigate. Doris even goes so far as to argue that it would be irresponsible to do otherwise.[6] As such, he thinks we should reject any conception of Aristotelian-styled broad or 'global' character traits and replace them with a more empirically accurate notion of narrow or 'local' situational traits, which are those that are relative to particular situations and are thus the morally relevant ones.

---

[6] Doris gives the example of meeting with a colleague with whom a subject has had a long flirtation while the subject's spouse is out of town. "If you are like one of Milgram's respondents, you might think that there is little cause for concern; you are, after all, a morally upright person, and a spot of claret never did anyone a bit of harm" (1998, p. 516). But since relying on such 'virtue,' that is, relying on something that empirically fails the test of explaining moral behavior will make it more likely that the subject will fall prey to infidelity than would be the case if he took the situational factors more seriously, "the way to achieve the ethically desirable result is to recognize that situational pressures may all too easily overwhelm character, and avoid the dangerous situation" (p. 517). As such, Doris claims that we have a "cognitive responsibility to attend, in our deliberations, to the determinative features of situations" (p. 518).

Along with Doris, Gilbert Harman (1999, 2000, 2009) has argued for the ineffectiveness and thus irrelevance of robustly global character traits for moral behavior. Whereas Doris claimed his ambitions were "modest" in that he merely wanted to "produce the beginnings of a suspicion that Aristotelian moral psychology may be more problematic than philosophers […] have thought" (Doris 1998, p. 505), Gilbert Harman goes all-out to argue that "it is better to abandon all thought and talk of character and virtue" (2000, p. 224). Given the kinds of experiments like 'coin-slot' and 'Milgram,' Harman claims that it may just be that "there is no such thing as character, no ordinary character traits of the sort people think there are, none of the usual moral virtues and vices" (1999, p. 316). As such, when we attribute the traditional kinds of character traits, we often make a huge moral error:

> In trying to characterize and explain a distinctive action, ordinary thinking tends to hypothesize a corresponding distinctive characteristic of the agent and tends to overlook the relevant details of the agent's perceived situation. Because of this tendency, folk psychology and more specifically folk morality are subject to […] 'the fundamental attribution error' (Harman 1999, p. 316).

Following Doris, Harman argues that the most relevant factors influencing behavior are situational ones. The relevant difference, then, between persons who act differently in a particular situation or in situations that are similar to one another, is not that some possess a certain trait and others do not. Rather, these persons "differ in their situations and in their perceptions of their situations" (1999, p. 329). As such, 'character' as we colloquially think of it has

no important bearing and is an erroneous way to both *evaluate* moral agents and *explain* why they behave as they do.

One interesting implication of this critique is how it extends to virtue responsibilist epistemology. On one hand, given that Zagzebski's epistemological theory attempts to parallel parts of Aristotle's moral theory, and the critique has been leveled specifically against the latter, it then ought to find a natural application to Zagzebski's theory. On the other hand, however, we should demand empirical evidence for this: do seemingly irrelevant features of situations actually play a more important role than intellectual virtues or character traits in whether or not a person behaves in an intellectually excellent way? Furthermore, do such situational factors actually play a more explanatory role in how people come to have justified beliefs and even knowledge, rather than their own truth-conducive intellectual efforts?

Mark Alfano thinks that they do. With Doris and Harman in hand, Alfano takes Zagzebski's conception of the efficacy of intellectual virtue to task. Zagzebski's claim was that being an *epistemically responsible* person (the epistemic counterpart of *morally responsible*) amounts to being motivated to pursue truth and thus motivated to develop and exercise *virtues* that successfully achieve that goal of truth. These become fixed in one's character and thus serve as the explanatory or causal *reason* why we get things epistemically correct when we do. Alfano, however, is curious whether intellectual virtues actually do this kind of work. If they do "then virtues had

better be explanatorily real" (Alfano 2012, p. 225). He wants to argue, however, that we have good reason to doubt this is the case since "most people's conative intellectual traits are not virtues because they are highly sensitive to seemingly trivial and epistemically irrelevant situational influences" (p. 234). Specifically, Alfano is interested in *epistemic* cases parallel to the public-phone-patron experiment noted above which examine "the influence of positive moods on cognitive motivation and processing" (p. 235).

Consider three empirical examples that he cites in order to set up his argument that responsibilist intellectual virtues, like their moral counterparts, do not play a relevant explanatory nor evaluative role in accounting for epistemic successes.

First, Alfano notes the Duncker candle task experiment[7] where subjects' "intellectual flexibility and creativity [were] tested" when they were given "a book of matches, a box of thumbtacks, and a candle," and subsequently asked to "fix the candle to a vertical cork board in such a way" that no wax drips when the candle is lit (Alfano 2012, p. 235). The interesting thing about this task, however, is that the experimenters designed it so that the only way to successfully complete it is to empty the box of its contents (the tacks) and use the empty box itself it as a candle-plate that will hold the candle and collect the wax as it melts from it.

[7] See Duncker (1945)

Second, Alfano cites experiments that required test subjects to perform the 'remote associates test' (RAT)[8] that presents them with various words that share a common conjoining word without revealing what that common conjoining word is. The idea, then, is to see how they figure out that associated word. In the particular experiment that Alfano cites, the experimenters[9] gave three words "'sore,' 'shoulder,' and 'sweat'" to subjects who were then asked to come up with the additional word that naturally conjoins to the three, which in this case was 'cold'[10] (Alfano 2012, p. 235).

Third, Alfano cites a study in which "researchers presented medical internists a task in which they had to identify the correct diagnosis of a hypothetical patient based on a description of his condition" (Alfano 2012, p. 237).[11] The question they were seeking to answer was given in order to study the kinds of external factors that helped the internists reach the correct decision in a faster and reliably accurate way.

What experimenters found in all of these cases is that "epistemically irrelevant mood elevators led to significantly increased performance" on all of these tasks" (Alfano 2012, p. 236). In the candle task experiment, subjects more reliably and more accurately reached the solution (the mere box itself must become the candle-plate) *only* if the tacks and the box they came in were

[8] See Mednick (1963)
[9] See Isen, Daubman, and Nowicki (1987)
[10] *Cold* sore, *cold* shoulder, *cold* sweat.
[11] See Estrada, Isen, and Young (1994)

presented as separate items. If they were presented together as a 'box of tacks,' then subjects failed at accomplishing this task most of the time. As Alfano argues, those that failed were not "intellectually limber enough to think of the box in a new way" (p. 235). In the RAT test, the subjects' performance was significantly improved when prior to the test "their mood was situationally elevated by showing them a short comedy film or giving them candy" (p. 236). Similarly for the medical internists, they were able to identify a "correct diagnosis" of the "hypothetical subject […] more quickly and with greater flexibility of thought" when they were given "a small gift of candy just prior to the experiment" (p. 237-8). So, whether it was the 'sugar rush' or just the mere act of receiving a nice treat, their mood elevation was the relevant factor in this and in the other two cases that caused them to perform epistemically successful.

For Alfano, this reveals two things that prove detrimental for responsibilist epistemology.

First, contra Zagzebski, the subjects who 'knew' how to successfully complete the tasks *cannot* be credited with reaching those solutions as a result from deeply entrenched intellectual virtue. Their success was not *due* to virtue, but was a result of situational factors: candy, comedy, and a varied presentation of the task at hand. Thus, their 'true beliefs' were acquired "not through flexibility and creativity" as global virtues but through flexibility *while in a good mood* and creativity *while in a good mood*" (Alfano 2012, p. 236; *emphasis* mine).

What does this mean, then, for a responsibilist theory such as Zagzebski's? Alfano argues that given's Zagzebski's account of knowledge (as well as justified belief) as being the result of an exercise of intellectual virtue instigated by the motivation for truth, she would have to commit to saying that these subjects didn't know the solutions they reached nor were they even justified in believing them since their beliefs didn't exemplify the exercise of intellectual virtue. Alfano claims, however, that this would be a problematic position to take because it is so counterintuitive: "it sounds more natural to say that they did know the solutions" even though the "funny video or candy lifted their spirits" and brought about the knowledge (Alfano 2012, p. 236).

Second, then, such findings spur Alfano to make the further generalized claim that

> many people do not possess creativity, flexibility, and curiosity as such, but inquire and reason creatively, flexibly, and curiously when their moods have been elevated by such seemingly trivial and epistemically irrelevant situational influences as candy, success at anagrams, and comedy films (Alfano 2012, p. 239).

Most people, then, "lack the consistent motivation required for intellectual virtue" when virtue is conceived in a 'responsibilist' way (p. 235). Yet, at the same time, people "know quite a bit" (p. 234).

What all of this suggests to Alfano is that "we should accept epistemic situationism and reject classical responsibilism" (Alfano 2012, p. 240). It is not obvious that responsibilist virtues do any epistemic work at all, but it does seem clear that seemingly epistemically-irrelevant features of situations is what

does the real epistemic work in them. In sum, then, Alfano notes that the situationist challenge to virtue responsibilism "can be framed as an inconsistent triad:"

> (non-skepticism) Most people know quite a lot.
> (classical responsibilism) Knowledge and true belief acquired and retained through responsibilist intellectual virtue.
> (epistemic situationism) Most people's conative intellectual traits are not virtues because they are highly sensitive to seemingly trivial and epistemically irrelevant situational influences. (Alfano 2012, p. 234)

5.4. One and the Same Problem

Could it be that the generality problem and the situationist critique are really, at root, different versions of the problem of getting at the right description that would satisfy our normative and explanatory needs and interests? There are several reasons to think this may be the case.

One reason may be because virtue-reliabilism and virtue-responsibilism *both depend on reliability as a crucial factor for evaluating epistemic actions*. As Guy Axtell (2010) has noted, even though the generality problem is usually "presented as an objection to reliabilist theories of justification, […] it is more accurately a problem that must concern any theory that has even a reliability component" (p. 74). For virtue-reliabilism, the degree of a belief-forming process's reliability is essential for conferring justification on the beliefs produced by that process. For virtue-responsibilism, how reliable the agent is in carrying out virtuous epistemic actions is an essential indicator as to whether

or not she actually possesses those robust traits that can "resist undermining" (Axtell 2010, p. 74).

Both the generality problem and the situationist critique similarly, however, challenge these reliability components by denying that reliability can actually be discovered or described precisely and accurately (for reliabilism), and in claiming that it doesn't really exist in a globally robust and cross-situational-consistent way since the traits it is said to be predicated of do not themselves exist in any meaningful obvious way (for responsibilism). As such, there are other salient factors for epistemic successes and failures rather than reliable belief-forming processes and reliably robust virtues, and those are the factors of situations themselves that are the relevant descriptors. What the *problem* and the *critique* poke holes at, then, is what it is that we may think is actually reliable.

This leads to a second reason why these problems are similar ones: *they gain their thrust by pointing out the saliency of the factors of situations in which agents act for how the epistemic process should be explained*. Just as both problems suggest that we cannot fully describe any such success or failure solely in terms of the agent's capacities or competences, they point this out via the need to cite some aspect (not too much or too little, but what is relevant to the successful working of the agent's belief-forming-process) of the situation in which the agent manifests her competences whether that be a Sosian competence or a Zagzebskian excellence.

Pollock & Cruz (1999, pp. 116-7) note this in their take on the generality problem, that to talk about a belief-forming-process's reliability (or, to put it in Sosian terms, an agent's belief-forming-virtue/competence) solely on its own merit *divorced from the circumstances* under which is works successfully and/or unsuccessfully is to do so deficiently.[12] So, for example, the process <*my vision-competence for identifying rare birds*> can be appraised as being reliable or unreliable only in reference to the circumstances under which it operates and thus performs, such as on clear days (where it is reliable) or on foggy ones (unreliable). As such, the generality problem is not a problem of specifying the epistemically relevant description of a mere mechanism (the reliability or unreliability of which determines justification or the lack thereof), but of specifying the epistemically relevant description of a belief-forming mechanism *relative to a certain set of external conditions.* (This is especially apparent in how Brandom framed the generality problem in terms of 'reference class,' which I noted above in §5.2.).

Similarly, Mark Alfano notes that the situationist critique is responsibilism's "own version of the [generality] problem" in that it challenges responsibilism to clearly define the level of generality to which virtues are to be defined: "should conative virtues be coarsely individuated, so that *open-*

[12] Hence, this is why I earlier referred to a belief-forming process as including not only the actual mechanism that works to produce the belief, but also the circumstances or conditions or situation (all synonyms here for my purposes) in which that mechanism works. As such, a belief-forming process = mechanism + conditions.

mindedness makes the cut, or should they be finely individuated so that *open-mindedness towards friends while in a good mood* makes the cut?" (2012, p. 233). As such, Alfano posits the situationist critique for responsibilist epistemology as 'the generality problem for responsibilism.' Its claim is that to understand virtue in a broad and coarse way is problematic due to what empirical research has told us about the inadequacy of such putative so-called virtues, namely, that they don't do the work that their proponents think they do. For Doris, then, the most promising way to think of virtues is as very *localized* instances of a trait where the trait is relativized to a particular situation (and hence relativized again to a different situation) instead of thinking of virtues as broad global traits. Thus,

> in the same way that the situationist challenge depends on whether moral virtues are coarsely individuated (e.g., honesty, courage, modesty) or finely individuated (e.g., honesty while watched by fellow parishioners, courage in the face of rifle fire, modesty before peers while in a good mood), the generality problem depends on whether intellectual virtues are coarsely individuated (e.g., curiosity, creativity, intellectual courage) or finely individuated (e.g., curiosity while in a good mood, creativity after being given candy, intellectual courage in the face of non-unanimous dissent). (Alfano 2012, p. 233).

As such, both problems challenge their target theories to provide a principled and empirically accurate way of describing the virtues that those theories rely on. So, these are reasons why the *problems* facing each theory are very similar.

I would also suggest, however, that not only does virtue reliabilism and responsibilism share these same problems, but they share in the possibility of

being interpreted as two respective prongs of the same project and *they are able to do so in virtue of their similar contextual structure*.

To illustrate this, consider the micro-contextualist dimension of Sosa's theory regarding attributing different *kinds* of knowledge. If the standards for knowledge, which have traditionally been held high by philosophers, also include the more relaxed object-level dimension, then we can be externalists *and* internalists regarding knowledge. We can say there are low-grade epistemic mechanisms (such as visual perception, memory recollection, etc.) that work without any need for the agent's conscious reflection and that yield true beliefs. Furthermore, we can also say there are higher-grade epistemic mechanisms or 'meta-processes' where a conscientiousness and an awareness of one's cognitive endeavors yield true beliefs of a higher, more sophisticated sort. And furthermore, it would seem that such meta-processes could confer justification on beliefs but only relative to a certain set of circumstances in a similar way that object-level processes do (e.g., the agent probably needs to be in good psychological shape, not drugged, not dehydrated to the point of 'seeing' mirages, etc.).

So, we can say that both kinds of contexts are conducive for producing justified beliefs, but we'll call successful instances of them at the object-level low-grade 'animal' knowledge, and successful instances of them at the reflective-level high-grade 'reflective' knowledge. What we find, however, is that Sosa's notion of 'reflective' knowledge is very comparable and compatible

to what Zagzebski has in mind when she argues for her analysis of knowledge that is motivated by a responsible, conscientious pursuit of virtue. As such, the two accounts may possibly describe two kinds of epistemic performances that we may achieve.

If this point is correct, and both accounts share in the possibility of a sort of harmony via their contextualism,[13] could this be helpful in thinking about a way to respond to the challenges posed by the generality problem and the situationist critique?

Although I will develop this more in the final chapter, it will be helpful to note here that for Sosa's virtue-reliabilism, an important feature that determines what kind of justification (either animal or reflective) is required for a particular kind of belief is the micro-context in which the agent is forming the belief. As I suggested earlier, for Sosa, it seems that the relevant degree of 'reliability' that will be required for conferring justification on beliefs made in particular micro-contexts will be mutually determined by the macro-context. If a particular micro-context's norms are such that it requires higher-

[13] Of course, one issue here is that while it seems that Sosa could concur with this sort of harmonizing, it's not clear that Zagzebski would. After all, as noted in the previous chapter, she sees Sosa's definition of 'virtue' as quite deficient in that, although it has Greek underpinnings, it's not as Greek as hers is – i.e., hers is more robustly (and specifically) aretaic. But one thing that her theory does suggest is that an intellectually virtuous person whose beliefs arise out such virtue will more often than not have some kind of reflective understanding or comprehension of that justifying factor – she might even argue that her account at least allows that potentially all beliefs out of intellectual virtue carry the possibility of being conscientiously reflected upon (except for babies and chimps, however, who are seemingly unable to have knowledge or, at the most, very little of it). As such, all justified beliefs for her might be what Sosa would call reflective beliefs. So, plausibly, while Sosa holds that there are such things as animal and reflective beliefs, for Zagzebski all justified beliefs are, in Sosa's terms, reflective ones.

order beliefs or lower-order ones, this is something that will be agreed upon at the macro level. This is because *we* determine the kind of 'reliability' that should be required for justification via our understanding of the particular kind of belief that's called for in that micro-context as well as our own needs and interests. As such, the sort of determination that's needed for an adequate and successful epistemic performance is a contextual one.

Similarly for responsibilism, the agent must figure out the best way to exercise her virtue, that is, the best way to manifest it so that her action doesn't go wrong in terms of excess or deficiency, but as the 'intermediate.' Similarly to Sosa's reliabilism, then, the criteria for this seems to be determined by the macro-context in terms of a mutual understanding and acceptance of what is the virtuous exemplar to be admired and imitated and thus exercised at that micro-contextual level.

And so, we notice the connection between the contextual nature of these theories and the challenges that have been posed. What the critics have been referring to as 'situations', are what I have been referring to as 'micro-contexts.' Doris (assuming he would agree with Alfano's analysis of virtue responsibilist epistemology) would argue that what we need to do is consider the situational factors in order to adequately describe and explain what influences people's epistemic actions and, thus, describe such traits at the micro level as relativized to particular situations.

The truth, however, is that this is something we've already been doing, both in our analysis of virtue responsibilism as well as in our analysis of virtue reliabilism. Micro-contexts have a huge role to play in understanding what the relevant virtues to be exercised in those contexts are because they must be contextualized to them in terms of being manifested in a particular way that is appropriate to the particular micro-context. This understanding, then, may point to a resolution to the challenges we have considered thus far in the chapter. I will attempt to develop such a solution in the next chapter.

5.5. Dissimilarity of the Problems

In suggesting that the challenges posed by the generality problem and the situationist critique are similar enough so that some kind of resolution to the former might also apply to the latter, there is at least one tension I should deal with here.

As already noted, virtue reliabilism holds that a belief-forming mechanism explains the *existence* of a true belief by being its cause, but as the generality problem points out, that mechanism cannot solely explain *why* it is reliable. For that explanation, we must accurately describe the circumstances in which it operates with such precision that we don't describe it too narrow or too broad. This is the gist of the generality problem. As such, virtues or processes have to operate in circumstances. They cannot operate outside of

them. However, the generality problem *does not* challenge the causal efficacy of the processes that form beliefs. It recognizes that they do indeed cause them.

The situationist critique, however, *does* challenge the causal efficacy of responsibilist virtue by arguing that *nothing about the agent herself* is either the prominent explanation *or the cause* of an agent's particular action, nor of why her action was right (or wrong). For those explanations, we must turn to the situation itself. So, as Doris puts it, 'global traits' do not explain; just local ones do. Virtue, thus, doesn't explain any action's cause nor does it explain the moral merit of action.

So, the analogy between the two problems is not perfect since they do indeed take a different perspective, not only on the content of virtue, but the active or positive (or empty) role of virtue. I do not think this is a major problem primarily because, at bottom, they both are seeking an explanation for epistemic justification, not necessarily for what *makes* a correct epistemic action (i.e., true belief) justified, but how we are to describe the sources of epistemic action that makes those actions justified. As such, both the generality problem and the situationist critique seek an explanation of how virtue explains why a justified belief is justified – not an explanation as to why the belief itself exists, but why it's justification obtains.

Chapter Six

A Contextual Resolution

> *There is no question of identifying the total cause of an agent's action or belief, and the partial causes we select and deem salient and usefully generalized upon, whether triggering or configuring causes, situational or agential, are contextual and have much to do with the interests-in-explanation of the persons providing the disposition-citing explanation.*
>
> - Guy Axtell (2010, p. 75)

6.1 The Attempt to Resolve the Problems in a Parallel Fashion

Consider again Mark Alfano's claim that the generality problem and the situationist critique "share many structural similarities" in that they both challenge their respective target theories (virtue-reliabilism and virtue-responsibilism) with the task of accurately describing the intellectual mechanisms (reliability's reliable competences and responsibilism's the deep-seated entrenched traits) that get us our justified beliefs as either "coarsely individuated […] or finely individuated" (2012, p. 233). What this means is that *we need to be able to come up with an adequate description* of the process that's actually doing the work in forming our justified beliefs so that we may accurately evaluate and appraise its product beliefs.

For virtue reliabilism and its conception of virtues – which for Sosa are those competences and abilities that allow agents to be successfully adroit – it's

not enough, however, to talk about the cognitive mere mechanism itself that forms beliefs. As I noted in chapter two, such mechanisms or 'virtues' never operate alone but always in some circumstance, always under some condition(s). Their analysis and appraisal must, then, always take place against a similar analysis and appraisal of the status of the conditions that they operated under. So, when we appraise a mechanism as being a reliable one, we mean that it is reliable *under x* conditions. And, so, those conditions under which the mechanism operates successfully most of the time are what we will call the *appropriate* conditions for its operation. As such, that mechanism *and* its appropriate conditions is a *reliable process*.

But as the generality problem points out, even this is not enough of a description. What is it about the conditions themselves under which the mechanism operates reliably which are the relevant factors affecting the reliability of the entire process? What aspects of the conditions should be included in the description of the process? Does the fact that I am wearing black socks instead of green ones make a difference as far as its reliability is concerned? What if it actually does, if in fact I am less reliable at forming justified beliefs if I wear black socks than if I wear green ones due to some psychological quirk of mine? How would I even go about delineating this particular non-epistemic issue as part of my belief-forming processes? So we see, the problem of generality is a problem indeed of describing the process,

but at its core is a problem of describing the conditions or circumstances that make the process a reliable one.

For responsibilist virtues a similar problem arises, as claimed by the situationist critique. We can't talk about responsibilist intellectual virtues as broad *global* virtues doing the kind of work that responsibilist theories have them doing because there is too much social-psychological evidence to suggest that situational features (and presumably those which have no epistemological relevance whatsoever) are the most salient ones which explain our epistemic successes. As such, the description of these kinds of virtuous processes need to be more inclusive, taking account of the *local* or *situational* aspects of an instance which an agent believes accurately. As such, it is unintelligible to construe responsibilist virtues as entrenched and consistent qualities of an agent. They only really exist (that is, if they exist at all) in light of the particular circumstances they are manifested in.

So, intellectual virtues such as *open-mindedness* and *intellectual humility* are not accurate descriptions of virtues. Rather, the description needs to be something like *open-mindedness-when-in-a-good-mood* and *intellectual-humility-when-in-the-presence-of-intellectual-superiors*, and so on. As such, the situationist critique requires of the descriptions of responsibilist virtues what the generality problem requires of descriptions of reliabilist virtues. That is, when we want to appraise others' epistemic performances as being successful and thus attribute virtue to them, we need to be able to offer an adequate

description of the conditions under which the agent-mechanism (i.e. the 'virtue') is said to have operated under.

That being said, given the similarity of the problems afflicting virtue reliabilism and virtue responsibilism, "a solution to the generality problem for responsibilism might also be exported as a response to the situationist challenge to virtue ethics" (Alfano 2012, p. 233). In response to Alfano here, I would like to take up his challenge and attempt some sort of resolution. In particular, I'm curious whether pointing out the contextual aspects of virtue reliabilism and responsibilism (as I have been doing throughout this dissertation thus far) might help in reshaping our understanding of the generality problem and take some of the sting away from its application. If so, then an inkling of a resolution might make itself available and applicable in response to the situationist critique.

6.2. A Contextual Attempt at Resolving the Generality Problem

Let's consider the generality problem. Given the social macro-contextual element of *reliability* in Sosa's virtue reliabilism, how does it affect the relevance of the generality problem for it? After all, the generality problem tells us that we must have a description of the belief-forming process that's not only precise but is the *relevant* one from which to gauge the level of reliability that confers justification on a belief. But if our epistemic desiderata (in particular, reliability and justification) are macro-contextually determined for

particular micro-contexts, does that in any way change how the generality problem applies to virtue reliabilism? In other words, does my nuanced contextual understanding modify in any way the applicability of the generality problem to Sosa's virtue reliabilism?

Mark Heller (1995) seems to think that a contextual understanding of how reliability works might dissolve the concerns raised by the generality problem since the reason it arises in the first place is because "we ignore the context-relativity of our use of the term 'knowledge'" in that "the [belief-forming] process has to be as reliable as the evaluator's standard's require" yet "those standards vary from context to context" (p. 501). Reliabilist 'standards' that prescribe what a belief must be in order to count as 'knowledge' more fundamentally, however, are going to require that the process that produced it meet some threshold of reliability in order confer justification on it. As such, these standards will be concerned with what 'reliability' itself means in terms of how reliable a belief forming mechanism must be in particular micro-contexts in order for its product belief to be justified as well as what are the particular situations or circumstances (or conditions) that that mechanism must operate under in order to even be able to be reliable.

In pointing out the context-relativity of epistemic standards, then, we find not only 'context-relativity' in how we attribute the term 'knowledge,' but something much more basic which is what allows us to context-sensitively attribute 'knowledge' in the first place: a context-relativity in how reliable a

'reliable' belief-forming process must be to produce a justified belief in one particular micro-context as well as in other micro-contexts.

As such, "reflecting the different concerns" we have "on different occasions of use, 'reliable' is context dependent in epistemological discourse" (Heller, p. 502). 'Reliability' is a "richly sensitive" term where an evaluator's "purposes, background beliefs, and standards" are going to play the major role in "positioning the 'enough' boundary" with regard to how much and what kind of reliability is appropriate for the micro-context in which it is discussed and applied and required (p. 504). Yet, in spite of the contextual complexity about how 'reliability' is used and referred to, this doesn't preclude our usual understanding of it, nor does it preclude its practical usefulness as a way in which we successfully structure and evaluate our epistemic activity.

Consider how contextual sensitivity affects different kinds of terms. For instance, Heller notes that while words like 'I' are sensitive to the context of utterance, they are so in a very straightforward way. For example, a major feature of the term 'I' is that its referent is always "the identity of the utterer," and so whenever the context varies, it will always do so "directly with that feature" (p. 503). As such, "it is not unreasonable to demand a general principle for determining its referent [that] will specify the relevant features of the context and their role in determining the referent of the term" (p. 503). But expressions such as 'similar enough', for example, vary with context in a way that is much more complex since "similarity judgments depend on a wide

variety of contextual features" such as "the evaluator's purposes, background beliefs, and standards"[1] since these factors "will affect her weighting of different respects of similarity as well as her positioning of that 'enough' boundary" (p. 504). Demanding some sort of principle that universally applies to every use of 'similar enough' in order to determine the validity and proper application of that expression would be quite inappropriate and probably absurd because given the nature of 'similar enough' it will never mean the same thing in every micro-context where it is applied. Similarly (no pun intended), if 'reliability' "is richly sensitive to the evaluator's context" in reflecting whatever the interests, needs, purposes, and perspective of the evaluator who is applying and attributing it, the generality problem's "demand for a [universal] principle evaluating every case" would "not [be] appropriate" (Heller 1995, p. 503).

So, then, is 'reliability' context sensitive in such a way that makes the generality problem null and void? Consider two of pieces of "evidence" that might help prove this.

First, any attribution of reliability will be done with certain standards in mind as to when a belief-forming process counts as being reliable; yet, those standards only apply to particular micro-contexts (or, at least, will not uniformly nor universally apply to all micro-contexts in the same way) and will

[1] Which, I might add, might be due to macro-contextual determination and pressure arising from the interests, needs, and purposes of the relevant epistemic community.

thus be context-relative. Why? Because, after all, at the macro-contextual level our needs, purposes, and interests as epistemic beings will seemingly influence and determine what the various epistemic norms and standards (including reliability) will look like for various situations. In this way, 'reliability' is doubly contextual: the proper threshold determining how much reliability is 'enough' reliability will be influenced by what the macro-level views as important for its interests, and then that gets applied in the form of norms and standards for the various individual micro-contexts given what's at stake in those particular situations. So, in one way, reliability will be 'macro' context-relative (i.e., relative to the particular macro-context and its needs/purposes/interests) as well as 'micro' context-relative (i.e., relative to the particular situation and hence not uniformly required across every situation or circumstance).

Heller recognizes that the generality problem involves the two important aspects of "describing the belief producing process" as well as "the problem of […] describing the environment in which the belief is formed," or, as I have been making this distinction, between the belief-forming mechanism (i.e. what the agents offers) and the conditions or environment under which that mechanism works[2] (p. 1995, 504). What Heller claims, however, is that recognizing this distinction and then focusing on the

[2] I have been making the distinction in my particular way so that I could refer to both the *mechanism* and the *conditions* together as the belief-forming *process* since that seems to be what the generality problem is requiring: an adequate description of the process that includes both the mechanism and the conditions under which it operates in order to gauge reliability. As such, to maintain a certain level of uniformity, I will continue to make this distinction in my way and present Heller with my particular terminology.

conditions/circumstances/environment prong of a belief-forming process makes the generality problem "sound a lot like the problem of relevant alternatives" (p. 504). And given that the recognition of the micro-context (i.e. the circumstances and their particular contingencies) plays the all-important role in how alternative possibilities gain their relevance (or irrelevance), this would seemingly also affect how we understand the generality problem as well since "the situation of the evaluator plays a significant role in selecting the relevant worlds" in which knowledge claims are to be evaluated (p. 505).

In chapter two, I argued that Sosa's virtue reliabilism is contextual since it makes an accommodation for two different kinds of belief justification: animal (object-level) and/or reflective (conscientious-level). What is interesting about these kinds of justification is how they can be applied differently and separately to various micro-contexts given norms of those micro-contexts. Such norms arise from the needs and interests of the macro-context (the epistemic community) for those micro-contexts at a particular place in time. But as Sosa explicitly noted, the community's 'need' here is the need to be able to 'mutually rely' upon each other as *reliable* epistemic interlocutors, participants, and community members. (And note how that this has affinities to what David Henderson referred to as the community's need for you to participate as an epistemic 'gate-keeper' and to what Robert Brandom referred to as playing the intersubjective role of an epistemic 'score-keeper,' which I discussed in chapter one). And this is a very interesting and

curious idea, because to be a 'reliable' epistemic community member is going to mean different things depending on the particular kinds of belief claims that are to be evaluated.

Sosa's framework, then, allows for this kind of relativity in one way by implicitly accommodating different *kinds* or *species* of reliability through positing the particular kinds of justification (animal and reflective) for different kinds of belief claims. And the reason why Sosa does this is because *we need both kinds of justification* in order to make sense of our everyday experience as epistemic practitioners since both kinds of justification serve our particular epistemic needs and interests. We need to be able to talk about (and subsequently use/attribute) positive epistemic status as a spectrum, from low-level stuff to the higher-level stuff. We need to be able to make sense of why and how we can be justified in believing one and the same proposition in one situation but not in another. After all, (hearkening back to my example from chapter one) my community may only require me to have object-level justification in my belief that I exist in the way that I think I do (my belief that AB^e) in the micro-context where I'm casting my ballot on election day. Yet, for my belief that AB^e to be justified in the same object-level way would be terribly insufficient in the micro-context of the philosophy class I sat in on (after having voted, of course) because the conditions there are such that it requires a more stringent kind of justification (i.e., reflective) that is more appropriate for the needs and purposes of the macro-context that are revealed in the

norms/standards of that micro-context. As such, in some micro-contexts, my belief that AB^e is reliably formed (and hence justified) and in other contexts it's not reliable enough (and hence unjustified), even though nothing about that particular belief that AB^e and the way in which it was produced has changed. All that has changed is the particular micro-context that the belief is being held in, and hence, the particular standards/norms that it is being evaluated against. In one context, its reliability was good enough. In another, it wasn't, and so it needed something a bit more.

Why is it that one particular micro-context requires more? Why is it that the threshold regarding how much reliability is required, for instance, may be higher in one micro-context than in another? Again, it depends on the purposes, interest, and needs of the evaluator. As discussed in chapter one, one way to think about this is in terms of the stakes involved. Maybe there's just more at stake in being able to philosophically justify my belief that AB^e as opposed to justifying it at my local polling place. And so, the kind of reliability that is required will be a different sort for each 'possible world' or 'relevant alternative situation' from which that belief claim would be evaluated.

Sosa's archer wonderfully illustrates this. Consider that the archery community will expect the expert archer's archery-competence-mechanism to be successful in some possible worlds but not in every possible world. They expect her to perform adroitly when the conditions are appropriate for archery: the wind is fairly calm and the air is relatively thin and clear and there is

nothing to obstruct either her view or the path of the arrow. And so, when they appraise her as being a reliable shot, such 'appropriate conditions' is the underlying reference class by which her arrow-shooting-competence is judged to be reliable. In appraising her as being a reliable shot, then, the community is by no means implying that she is reliably successful at shots on Mars, or the Moon, or under water, or in dark caves, or while hallucinatically drugged. Nor would her unreliability under these kinds of conditions mean that she is unreliable *tout court*. To the competitive archery community, possessing this *kind of reliability* would be irrelevant and inappropriate for its particular needs and purposes. Yet, even if she could successfully perform in those particular micro-contexts, and if she was thus reliable in them, no doubt, that kind of reliability would be much different than the earthly kind.

In a similar manner, then, our macro epistemic needs and interests do not require of epistemic agents that they be 'reliable' in ways that are irrelevant or inappropriate to the normal pace and capacity of human cognition and practical life. And so, appraising an agent as a 'reliable' epistemic agent implicitly means 'reliable in her belief claims in normal human micro-contexts.' Of course, if an agent could indeed reliably form true beliefs about super-human things in a super-human way, then she would be reliable relative to those conditions as well. But this isn't necessary in order to be *reliable* in the way we usually mean it. Furthermore, we don't even require one to always

'know' things in a fully human or *reflective* way. If someone has a mere *animally*-justified true belief, for some micro-contexts that's good enough.

The kinds of epistemic abilities that we want each other as epistemic practitioners to have are those "aptitudes" that are always "relative to certain background conditions" which are "distinctive correlated parameters" determined by the needs and interests of the macro-context (Sosa 2007, p. 83). Such needs and interests are those that are the normal and practical ones manifested in our everyday epistemic tasks. One's attempts that fail "in abnormal circumstances do not show lack of the ability" and thus do not automatically make that person unreliable or undependable *tout court* (Sosa 2007, p. 84). "Despite such failures we might still depend on you in normal circumstances. What is required is only that your attempts tend to succeed when circumstances are normal" where such a "normality" is "determined implicitly […] by those who share the concept" (Sosa 2007, p. 84).

As such, the macro-context – playing the role of the belief evaluator – which sets the norms/rules for our epistemic practices including what are the conditions that are appropriate for our epistemic exercises, share in some sort of recognition of what that all means merely through the way it *practices* epistemically. As such, there can be implicit agreement on what it means to be a 'reliable' epistemic agent in a medical laboratory and a 'reliable' epistemic agent on a baseball diamond even though satisfying the reliability threshold will be very different for each particular practice. So, even though 'reliable' will

mean different things in different micro-contexts, successfully pointing out meaningful referents of 'reliability' isn't a mystery since micro-contexts themselves contain the standards that sort this out. Analogously,

> In a typical context in which someone says 'the refrigerator is empty' or 'the cat is in the box,' reference is somehow successful in spite of the fact that there are many refrigerators, many cats, and many boxes in the world. If we think of the problem of generality as a species of this problem it seems even more plausible that the solution should come from a proper appreciation of the role of the evaluator's context in epistemological discourse. (Heller 1995, p. 505).

Secondly, then, given that the meaning and application of 'reliability' is affected by the conditions within particular micro-contexts themselves, this reveals the 'richly sensitive' nature of 'reliability' in that it is affected by micro-contextual pressure. Because of this, it thus allows for shifts in meaning. It is inappropriate (and unrealistic) that there should be some principle or rule that universally sets a standard for how to describe 'reliability' generally across all micro-contexts: "the demand for a principle for evaluating every case is not appropriate for terms that are richly sensitive to the evaluator's context" (Heller 1995, p. 503).

Consider Heller's first car, "a then 15-year-old VW Beetle" that would sometimes "start with the turn of the key, but usually [he] had to push-start it" (p. 503). Now he owns a "Honda Accord, purchased new two years ago" which "has never failed [him]" (p. 503). For Heller, "the difference between the Accord when it starts and the Beetle when it started is that the Beetle's success was just a matter of luck […] because the process by which this car

started [was] not reliable" whereas "the process by which [the Accord] started [was] a reliable one" (1995, p. 503).

> If there were a real problem of generality for the reliabilist's explanation of the difference between knowledge and mere true belief there should be just as much of a problem for our present explanation of the difference between the two cars. Since there is no problem with the latter, there is also no problem with the former. Just as I can use the everyday term 'reliable' when describing my car, the epistemologist can use that same term in the same way when describing belief-producing processes. And just as one can unproblematically assert that a car is worth buying only if it is reliable, an epistemologist can unproblematically assert that a belief is knowledge only if it is produced by a reliable process. These considerations suggest that the problem of generality is a red herring. There is no real problem. There only seems to be a problem when reliabilists mistakenly accept the challenge of producing a principle for determining when someone has satisfied the reliability condition (Heller 1995, p. 503).

Consider that what made the Beetle 'unreliable' was simply a bad electrical connection between one of the spark plugs and the distributor, and so "a little jiggle" of the wire before attempting to crank it each time did the trick that allowed it to start (Heller 1995, p. 508). But notice what this means: it was *unreliable* under conditions where the wire wasn't checked and *reliable* under conditions where it was. And, so, Heller's complaint that his 'unreliable' Beetle only started 'with luck' instigates his friend's response that, "given that the wire was in fact snug at [the] time, what else could the car have done other than start?" (p. 508). Heller, then, is forced to agree with this: "the process that was used in this instance was the completely reliable process that includes the wire's being firmly in place" (p. 508). But at the same time, this does not imply that Heller was wrong or misguided in his original complaint about his Beetle's 'unreliability'. It simply means that in this conversation with his

friend, the context was shifted when the friend "change[d] the standards for process re-identification" by talking about the Beetle's reliability relative to a certain set of conditions (p. 508). "By emphasizing one particular feature of the process – the snugness of the wire – she made that feature most salient" and therefore "at that point" made it so that "the appropriate way to characterize the process would be to include that feature" (p. 508).

Similarly, Heller's original claim that the Accord is a reliable car is also subject to contextual pressures when, for example, that same friend points out that the Accord's spark plug wire is only a "big jiggle away from being loose" (Heller 1995, p. 509). Here, the friend is describing conditions that would make the Accord unreliable, and as such, "is raising the standards for reliability" within one and the same conversation and is "forcing into the domain of discourse worlds that had previously been irrelevant" (p. 509). Of course, Heller being the "accommodating conversational partner" that he is, "do[es] not challenge her new standards but instead generalize[s] [his] description of the process type so that it can occur in the newly relevant worlds" (p. 509). As a result, accommodating his conversational partner and re-identifying the process of starting the Accord *now* makes it 'unreliable,' but this is only because a new micro-context has been introduced that is the backdrop of this discussion of the Accord which is *now* taking place. Re-identifying the process is now necessary given the new possible world – the new set of conditions – that we must take into account in order to truly

determine whether or not the mechanism that is supposed to start the Accord could reliably do so under those conditions. The result we come to is that it could not reliably start the Accord under those conditions, but again, this wouldn't make the Accord unreliable *tout court* but *only in this new possible world of a 'big jiggle' that could happen*.

For Heller, this seems to show that as the "line" between "realistic and non-realistic possibilities, between close enough and not so, fluctuates with contextual factors, we should expect that in general the positions of the line can change even within a single conversation" (Heller 1995, p. 508). This line, then, "can be pushed further out as our standards for reliability increase, or it can be pushed further in as our standards for process re-identification increase" (p. 508). But just because such a shift can and does happen (and thus, the Beetle goes from being 'unreliable' to 'reliable' under a particular set of norms for what counts as being reliable, and the Accord goes from being 'reliable' to 'unreliable' under a different set of norms for reliability), this "does nothing to detract from the truth of [our] original claims of reliability" but "just shows that our *use* of 'reliable' is affected by contextual pressures the way we would expect of a term that is richly sensitive to context" (p. 509).

And so, for the generality problem to "demand a *fixed* principle for selecting the correct level of generality" for our epistemic claims is "unreasonable" since "what counts as correct varies from [micro] context to [micro] context" in that our particular epistemic norms that determine the

proper threshold (and kind) of reliability vary from micro-context to micro-context (p. 503).

There is something that seems intuitively correct about this sort of analysis. The reliabilist is not merely asking for an impossible principle of a 'correct level of generality' regarding a process type in order to gauge its reliability and hence the justificatory status of its resulting beliefs, but is *asking for the unnecessary*. We seem to operate with some kind of pragmatic guideline as to which level of reliability is appropriate for the various epistemic micro-contexts we claim true beliefs in and we seamlessly operate from it given that when those in our epistemic community challenge it, we adjust our descriptions. We re-identify and re-describe our belief claims in light of the particular micro-contextual perspective from which those challenges and requests for clarifications come.

Of course, this may be a bit optimistic. After all, we don't always agree as to whether or not our belief claims are justified, much less cases of knowledge. But this is merely an aspect of what it means to be epistemic practitioners in general, and belief-appraisers/evaluators specifically. The salient features of conditions for belief production and possession will be different to a certain extent for each epistemic agent. After all, Heller was complaining about his 'unreliable' Beetle to his friend who disagreed with Heller's particular assessment because for her the Beetle was quite reliable. However, both assessments were true and justified in that they each arose from

different yet reliable micro-contextual perspectives, that is, from different notions of what the relevant circumstances were *and* reflected the norms for how reliability obtains respective to those circumstances. Heller was right that the Beetle was 'unreliable' but only in terms of a particular micro-context: *before* checking the spark plug wire connections. And his friend was also simultaneously correct that the Beetle was 'reliable' but, again, in terms of a *different* micro-context: *after* checking the spark plug wire connections. The relevant difference that makes a difference in attributions of reliability, then, are differences in the micro-contexts from which the *evaluators* make their appraisals.

Consider what it means, for instance, to be a 'reliable' or 'unreliable' fisherman with 352 catches per year on his tally. If the micro-context from which I make my evaluation is one where he is the sole beneficiary of his fishing practice, then I would say he is quite reliable. After all, he's keeping himself fed on a daily basis. But if the micro-context is such that the local cannery is depending on him to feed the entire town, I would say he is quite unreliable at that. These opposite and contrary evaluations and appraisals can be made about this very same person without changing any feature about him or about his total lifetime harvest. But when the micro-context shifts, so does his reliability rating.

In light of this, consider my belief I discussed earlier regarding my own existence (in the way that I think I exist) that AB^e and how nothing about it

changes from one micro-context to the next. But if it inherits its justification from the process that produced it (as reliabilists claim), then the standards for reliability in the micro-context of the polling place are such that its generative process is 'reliable', yet at the very same time is 'unreliable' in the micro-context of the philosophy course. All that has changed, mind you, is the evaluator's perspective. She has examined the belief-forming process's reliability from one perspective and got a determination of reliability there, and also examined it from another perspective and got a completely different determination of reliability there.

For Heller, the generality problem

> is not built around any difficulty in specifying the subject's overall reliability rating. It is, rather, centered on the problem of specifying which overall reliability ratings count as knowledge – which are to be included in the extension of "knowledge". What the evaluator's situation determines is which sort of connections are relevant – which situations the evaluator cares about in the context in which she finds herself. […] Once we recognize the role that context plays in determining the truth of knowledge claims, we should see that there is no problem of generality. There is the fact that there is no principle that determines an appropriate level of generality for every case. (Heller 1995, pp. 513-4)

The conclusion argued for here, then, is that reliability is a standard of epistemic justification that is context-sensitive. After all, for some micro-contexts, such as one's grade 5 science project, a bare minimum threshold of reliability (say, 51% reliable) may be good enough to justify one's belief that 'the microscope is in good working order,' while in other micro-contexts, such

as the testing stage of a new vaccine, a much higher threshold is required by the macro-context.

Furthermore, consider Norman the clairvoyant (from §2.2). As noted earlier, people have different intuitions about this case. Some think that he exercised some sort of basic epistemic reliability given his clairvoyance and thus his beliefs formed from that competence are justified, but barely. Others, however, think that because he doesn't (and really cannot) have a reflective meta-perspective of his cognitive operation (and thus cannot examine his own belief-producing practices from this perspective), his clairvoyant beliefs are not justified and thus saying whether or not they're reliable is irrelevant. If Heller is right, the disagreement between these two positions arises out of a comparison of non-comparables – like comparing apples to oranges – because both evaluators have different conditions in mind. Truth is, they're both correct. Norman is a 'reliable' believer if the circumstances-prong of his epistemic situation is understood as a micro-context whose norms allow for his clairvoyant capacity on its own to be 'good enough', and he is an 'unreliable' believer in micro-contexts where the standards require him to have a cognitive self-reflection on his epistemic abilities.

And, so, given what the context at hand *is* in terms of the norms it imposes on epistemic practitioners like Norman, this will tell us whether or not Norman indeed is justified in his beliefs in that micro-context. For instance, we might illustrate this such as if he were an employee of a major

news organization. But to such an organization, Norman would be both an asset and a liability. He has this remarkable ability to instantly come up with accurate beliefs. However, he has no way to verbally justify them. Hence, given the (hopefully high) epistemic standards that are imposed on those who report the news, this is a role that Norman could probably never do. This is an epistemic practice that he would be unreliable at since a major part of this job will be the task of coming up with, evaluating, criticizing, citing, and delivering *reasons* – something Norman cannot do, but is something demanded by that particular micro-context. However, he would be extremely reliable in a position where he could inform reporters of his beliefs in order to get them a place to start investigating their stories. Indeed, he might be even more reliable than the best police scanners on the market. Here, then, is a context at hand, characterized by its norms (but even more fundamentally by our needs and interests that shape those norms), where Norman is reliable and justified.

One question that arises is whether or not we can legitimately do this sort of analysis with a concept like reliability. That we can work with such different uses of 'reliability,' even though those uses retain a similar nature or essence of meaning, may simply be an extension of what David Kaplan (2008) has argued regarding names and demonstratives, that with such terms there are *two kinds of meaning*. First, there is the *character* of a term and it is the "sense of an expression which determines how the content is determined by the context" (p. 359). As such, it is the rule we use for attaching a particular

referent to the term, that is, it 'picks' out the content of a term. And, secondly, there is the *content* of a term which "is always taken with respect to a given context of use" (Kaplan 2008, p. 358). The content, then, is the actual referent of the term and can shift and change depending on the context in which it is uttered or applied.

Applying this distinction to 'reliability,' we get some inkling of how it can specifically mean (content) different things in different micro-contexts because there is a general meaning (character) that is the rule that competent speakers of the language apply in those different micro-contexts where 'reliability' is employed. This 'character' aspect, then, explains why we are sensitive to such shifts in the meaning (content) of reliability (that is, its varying levels of degree) with seamless precision.

This may help explain why the level of reliability for the self-reliant fisherman is insufficient for the local cannery, why Heller's Beetle is 'unreliable' to him and 'reliable' to his friend, and why Norman is both 'unreliable' and 'reliable' at the same time: it all depends on how the *content* of 'reliability' has been determined by the *character* use of the macro-context (i.e., the particular way in which the 'rule' is being applied) for the particular micro-context from which our evaluations have been made. This requires us, however, to make the further claim that correctly applying the rule (or character) in a particular context requires that we have learned from our community of language users how to properly apply the rule for each context

that will then determine the content of the term. Thus, learning the proper character of the concept of reliability relies on our participation in the epistemic/linguistic community we are part of.

Thus, our use of epistemic language and epistemic concepts, be it knowledge attributions such as Sosa's micro-contextual categories of low-grade animal knowledge or higher-grade reflective knowledge, or justification/reliability claims that can easily shift either in a more strict or loose way depending on the context in which it is applied, all heavily rely on contextual features. And at the top of the list of such features is the macro-contextual aspect where we as the epistemic community learn how to apply and employ various usages of these concepts as well as collectively govern their proper and appropriate usages.

So, this putative resolution to the generality problem doesn't solve the generality problem *per se* on the latter's own terms, but makes an argument as to why it is inappropriate and inapplicable. The generality problem challenges reliabilism to provide a non-ad-hoc rule that guides and guarantees the *proper* description of a belief-producing process: one that accurately accounts for the level of reliability that confers justification on its product beliefs. The contextual response, however, claims that given the context-relativity of our knowledge claims, belief evaluations, and appraisals, this is impossible, and even more so, unnecessary. It's not a problem that different descriptions of the process get us different results. Indeed, when we evaluate our belief-forming

mechanisms in light of different circumstances or 'micro-contexts' and thus get different results as to that mechanism's reliability, this is a good and desirable feature of our epistemic practice in that it allows us to be more precise and exhaustive regarding our epistemic claims rather than blanket-evaluating our mechanisms as *tout court* 'reliable' or *tout-court* 'unreliable.'

Furthermore, this sort of contextual resolution is something already implicit in the virtue-reliabilism of Ernest Sosa, but it just needed to be developed a bit more explicitly, which is what I have tried to do. His distinction between animal and reflective justification mirrors different kinds of belief-process reliability, the positing of which just merely reflects our needs, purposes, and interests as an epistemic community or macro-context. As such, our various descriptions of a belief-forming process will get us the various results of 'reliable' or 'unreliable' but not in an ad-hoc way. Rather, those various appraisals are rule-governed in the sense that they are always made in terms of a particular micro-context, that is, for a specific situation or circumstance that contains its own particular standards for what counts as 'reliable' or 'unreliable' in that particular micro-context.

To answer Brandom's question from §2.2 – 'which reference class is the correct one?' – we simply answer: 'it depends which situational perspective you might be evaluating from' and thus 'it depends what your own needs, purposes, and interests are.' Far from being a hindrance to our epistemic practices, this kind of sophistication and complexity is quite a remarkable demonstration of

our epistemic precision. As such, there cannot be any rule that allows us to describe the one-and-only correct level of generality. We need to be able to talk about processes in different ways, all micro-contextually relevant, which serve our macro-interests and macro-needs in those micro-contexts.

6.3. A Contextual Attempt at Resolving the Situationist Critique

In short, my criticism of the generality problem is that despite its effort at precision regarding the processes that form our justified beliefs, it ignores the various macro (social) and micro (situational) features of our epistemic practice. There is no single level of generality to identify in order to gauge for belief justification. Depending what our epistemic needs and interests are, and hence depending from which sort of micro-context our evaluation and appraisal proceeds from, various levels of generality may serve our purposes at that moment and this sort of variation may present itself in the various sorts of conditions or circumstances from which we play our role as evaluators. The contextual sensitive nature of our usage of 'reliability,' both *implicitly* in the sense that we 'mutually rely' on one another to be good and faithful epistemic partners, and *explicitly* when we make belief evaluations and appraisals, suggests that we do indeed put it to use in many different ways given our need, interests, and subsequent norms for various micro-contexts.

Given that the situationist critique – the problem of how we should describe character traits as either more narrow and local or broad and global –

is similar to the generality problem, does a similar sort of response to what I presented in the previous section present itself here? Is it as simple as saying that for responsibilism, whether or not someone has acted out of virtue (and is thus 'virtuous') or not is merely a matter of what the evaluator's contextually salient needs and purposes and interests are at the time of the evaluation? But if we go along with this sort of explanation, wouldn't this seem to fly in the face of the aretaic notion that virtues are dispositions of an agent? Wouldn't this change the traditional aretaic understanding of virtue so much so that it'd be unrecognizable? Two issues arise from these questions that we must address and will then lay the groundwork for an explanation that is parallel to the one argued earlier in response to the generality problem.

First, there is the question of *how* situational or 'micro-contextual' we can describe responsibilist virtues while still retaining some semblance of their nature as *dispositional* traits rather than mere *situational* ones.

After setting out his situationist critique, Doris responds with his own remedy: reject all conceptions of Aristotelian-styled broad or 'global' character traits and replace them with what he considers to be a more empirically accurate description of character traits as narrow or situational traits relative to particular situations. As Alfano notes, his "cutting the fabric of traits so finely" through his "principle of individuation" of particular traits relative to particular situations" allows Doris to make extremely specific attributions of character, even to the point of not only differentiating *courage* as "*battlefield physical*

courage, *storms physical courage*, *heights physical courage*, *wild animals physical courage*" and so on, for instance, but "he even seems willing to individuate between *battlefield physical courage in the face of rifle fire* and *battlefield physical courage in the face of artillery fire*" (Alfano 2012, p. 231).

In terms of responsibilist intellectual virtues, then, those such as *open-mindedness* and *intellectual humility* are described too simplistic and broad and thus are not good and accurate descriptions of these traits. Rather, the description needs to be something like *open-mindedness-when-in-a-good-mood* and *intellectual-humility-when-in-the-presence-of-intellectual-experts*, and so on. As such, evaluations of intellectual character that make broad appraisals in the traditional sense should be scrapped and replaced with the situational features that play a role in what is actually *seen* or *experienced* by the evaluator in such a way that allows her to make her evaluation (i.e., trait ascription).

And second, there is the related issue, then, of whether or not these two sorts of descriptions of virtues – broad dispositions or situational traits – are mutually exclusive, and hence incompatible. This is Doris's claim. As Guy Axtell notes,

> Doris holds that one of the key upshots of the social psychological research is a "'fragmented' conception of character which countenances a plethora of situation-specific 'narrow' or 'local' traits" that aren't unified with other traits. […] Doris's fragmentary account holds "that systematically observed behavior, rather than suggesting evaluatively consistent structures, suggests instead fragmented personality structures – evaluatively inconsistent associations of large numbers of local traits." (Axtell 2010, p. 82)

As such, Doris posits that virtues aren't connected in any way, even if we talk about, for instance, the intellectual virtue of *intellectual-humility* and *intellectual-humility-when-in-the-presence-of-intellectual-experts*. For Doris, the former description of virtue as dispositions makes no sense (indeed, for Harman, it doesn't exist in any meaningful way). What we get, then, are a mishmash of localized traits with no essential connection to one another. They are merely context specific in a very strict and rigid way regarding an agent's behavior and not her disposition(s).

In response to these two issues, that of (i) maintaining a dispositional element in responsibilist virtue and that of (ii) the mutual exclusivity of dispositional and situational descriptions of traits, I'd like to suggest that (i) can indeed be accomplished while at the same time accounting for the local or situational aspects of trait exercise and attribution, and hence also suggest that (ii) is misguided. Indeed, the exercise and subsequent attribution of responsibilist virtue necessarily involves both dispositional and situational elements.

As such, my response to these concerns is motivated by an apprehension I have regarding the adequacy of a mere solely 'local' way of thinking about character traits. When we consider the moral theory made by Aristotle himself and from which we get a very robust explanation about virtues as dispositional traits, we are forced to notice that he made a very clear and important 'situational' point in his theory (which I discussed at length in

chapter three) that even though a virtue – this deeply entrenched dispositional trait – is exercised by the virtuous person in every situation that calls for it, the particular way in which it gets *manifested* varies from situation to situation. Hence, Aristotle presents virtue (whether a passion or action) as being an intermediate state that strikes a 'mean' between two different ways of going wrong,[3] and this gets fleshed out when the agent acts in particular situations. As such, there is already a sense in Aristotle's analysis in which virtues *are* localized. Although virtues are understood to be *dispositionally global*, they are always *locally manifested*. And the agent's ability to be context-sensitive in how she goes about determining the best way to manifest a particular virtue in different ways in different situations is itself, then, a result of the agent possessing virtuous dispositions.

Now, consider Zagzebski's analysis of intellectual virtue which is relevantly parallel to Aristotle's analysis of moral virtue. There are situations where the manifestation of the intellectual virtue of open-mindedness, for instance, will look one way in one situation and a different way in a different situation while for each situation or micro-context it is motivated by my desire for truth.

So, for example, the particular manifestation of my truth-conducive virtue of *open-mindedness* may be 'more open' in a situation where I'm

[3] "Virtue, then, is a state of character concerned with choice, lying in a mean, i.e. the mean relative to us, this being determined by a rational principle, and by that principle by which the man of practical wisdom would determine it" (1106b36-1107a5).

surrounded by intellectual superiors such as the world's top hepatologists (let's call this situation SIT-HELP) who are discussing my particular liver condition with me and prescribing the actions that I need to take in order to maintain the health of my liver. And, yet, that very same *open-mindedness*, however, may be manifested in a way that is 'less open' or 'more guarded' in a situation, for instance, when I get an email from the U.K. (let's call this situation SIT-SCAM) informing me that my 17th cousin died and left me £100,000,000 and all that I have to do to claim it is send my current address, birth date, and social insurance number to some guy in Manchester. As such, not only will the manifestation of open-mindedness look one way in SIT-SCAM and another way in SIT-HELP, there is a sense in which *it ought to* since this is exactly what would be expected of me from the *macro*-context, which I argued, is embodied in the notion of the virtuous exemplar. To intellectually behave well in these ways is to do as the exemplar would do in similar situations. And, so, the dispositional element of this virtue and its local aspects all work hand in hand. That I have that particular disposition is what allows me to apply it (and hence manifest it) in different ways given the particulars of the situation or micro-context that calls for it.

To illustrate this even further, consider again my use of David Kaplan's (2008) theory that names and demonstratives carry *two kinds of meaning*. Whereas the *character* of a term is the general rule we use for 'picking out' the content of a term, the *content* itself is then determined with "respect to a given

context of use" (p. 358). The content, then, is the actual referent of the term, which will look a particular way depending on the context in which it is applied, and then look a completely different way in another different context. If *virtue* (in the responsibilist way) can be thought of under this kind of explanatory framework, then we have some inkling as to how virtue can be both dispositional and situational all at the same time.

For instance, expressing the intellectual virtue of open-mindedness in SIT-HELP to a greater degree than I do in SIT-SCAM doesn't foreclose the fact that I am indeed expressing open-mindedness in both cases, even though their particular manifestations look very different. How can I do this? Because I possess this dispositional trait, this virtue, and in possessing it dispositionally I possess the (Kaplanian) *rule* or *character* of that virtue since I can apply that rule in certain ways given specific micro-context in which I'm using that rule. After all, this might be evidenced by whatever similarity or affinity that the manifestations of open-mindedness in SIT-SCAM and SIT-HELP have to each other (even though they are different by a wide degree, yet nonetheless there is an inkling of similarity). And, so, because I possess this virtue – and hence I possess this rule or basic character understanding of that virtue and how to apply that virtue – given the situation where I need to exercise it, I can weigh the contextual needs and pressures and then manifest that virtue with a specific content tailor made for that situation.

The 'character' (i.e. dispositional) aspect, then, of virtues allows us to apply them in specific situations (micro-contexts) where such manifestations are the specific micro-context-appropriate 'content' of the virtue. And, of course, this makes complete sense. It's not as though virtue is either dispositional or situational. Rather, because it's dispositional (has a certain *character*), and because we have *phronësis*, that dispositional trait, then, can be applied (i.e., it can be given specific *content*) by us in various situations.

And, so, given the contextual aspect of virtue, that is, of the agent's ability to manifest virtue in a variety of ways, each way being an equal expression of the virtue in question, a contextual response to the situationist critique may be achievable. And this may be accomplished while maintaining the dispositional aspect of responsibilist virtue *and* accounting for the important influence that aspects of the situations or micro-contexts themselves have on the way in which we manifest virtue.

Kristján Kristjánsson, for instance, has argued that

> Even if social psychologists succeed in convincing us that people do not possess character traits qua robust behavioral dispositions, it does not mean that people do not possess character in the more nuanced Aristotelian sense, which is holistic and inclusive of judgment, emotion and manner, as well as action. (Kristjánsson 2008, p. 67)

After all, when we consider all of the contextual factors that influence the exercise and subsequent appraisal of virtue, there is much more going on besides the mere 'behavior' or 'action' that is seen on the surface. As Kristjánsson notes further,

> The typical Aristotelian virtue will nevertheless be a complex character state (*hexis*), at the core of which lies moral sensitivity, exhibited through emotional reactions, to goings-on in the world around the agent. Whether a felt emotion should be acted upon or not is always a separate question, however. The answer to that question must take various situational factors into account and be adjudicated through the intellectual virtue of *phronesis*. The description of a given virtue is not fully exhausted by characterizing its underlying emotional sensitivity and the range of possible actions to which it can give rise. Virtuous persons also comport themselves in certain distinctive ways which reverberate through all their attitudes and conduct; what matters is not only what they feel and do but also the manner in which they feel and do it. Great-minded or magnanimous persons, for instance, exude an aura of proper dignity and moral superiority which distinguishes them from other virtuous (but non-magnanimous) persons. This is why each *hexis* is truly a complex state, rather than a mere disposition to feel and act. (Kristjánsson 2008, p. 73)

In a similar way to what was noted in the previous section, then, there may be a need for a broad description and attribution of virtuous processes (reliabilism) and dispositions (responsibilism), or, the need may be for more narrow and fine-grained ones. And so, in Heller's words, we 'move the line' from narrow to broad, or from broad to narrow, given what our purposes in trait ascription may be. Maybe it's the case that we're faced with a number of situations that call upon the same kind of virtue (even though it will be manifested in each of those situations in a different manner), and so broad trait ascription is a helpful way to *generalize* or re-identify the trait in such a way that will be useful for our particular needs. As such, the situationist critique proves neither the non-existence nor meaninglessness of trait ascriptions in a broad dispositional way. But what this has done is force us to rethink how virtues like this may work, and such an analysis actually goes so far to show us how indispensible they may be.

Conclusion

So what have we ultimately concluded? Well, for virtue reliabilism, what makes it a contextual theory is the fact that its notion of justification is itself contextual: a macro-context determines what counts as belief-justification for the various situations we find ourselves in. And since justification is understood in terms of reliability, what it requires is a proper or sufficient threshold of reliability to be prescribed for certain situations. What determines the suitable level of reliability, however, is our needs, interests, purposes, etc. – i.e., macro-contextual 'stuff.' The generality problem seemingly assumes that this 'stuff' is unimportant given its demand for us to come up with a principle or set of rules by which all processes – that is, our cognitive mechanisms and the appropriate conditions for their operation – *must* be described in order to objectively and accurately determine their reliability.

But as I've tried to show, where the generality problem may go wrong is in its assumption that being rule-governed necessitates playing by one set of rules. What experience shows us, rather, is that there are all sort of micro-contexts governed by the rules and norms of the macro-context (whether broadly construed in terms of merely being a member of the human epistemic

community, or more narrowly construed such as being a member of a specific specialized community [i.e., medicine, biology, chemistry, religion, education, etc.]) whereas such rules and norms are geared towards articulating what it is to have justified beliefs in a particular micro-context.

For virtue responsibilism, then, what makes it contextual is the fact that we can distinguish what it is to possess a virtue from what it is to actually demonstrate or manifest that same virtue. It makes complete sense to say that one can possess the truth-conducive epistemic virtue of open-mindedness, for instance, across a broad range of circumstances and situations (micro-contexts). But to be a responsible epistemic practitioner means that we have to figure out how best to manifest that virtue in various micro-contexts. Of course, we're not without help. The standard bearing *exemplar*, which is macro-contextually determined, is supposed to help in this role.

The situationist critique claims that because virtue is merely situational (given the results of social psychological experiments), virtues are not possessed in any robust deep-seated way. But what I have tried to argue is that Aristotelian-styled aretaic theories, especially virtue responsibilism, have a contextual nature that recognizes both dispositional and situational aspects. Indeed, I have tried to do some tedious hermeneutical work in Aristotle, MacIntyre, and Zagzebski to show that they seem to imply this themselves, although not in the explicit way that I have tried to explain it. The result is that although virtue is understood as being dispositional, the situational

demands of a micro-context will (and ought to) influence and guide how the agent manifests a particular virtue in that micro-context.

For both virtue-reliabilism and virtue-responsibilism, then, we can happily hold that belief justification has a contextual, albeit not relativistic, nature if, as Kaplan's analysis suggests, certain kinds of terms and concepts have two specific kinds of meaning. But this obtains if *reliability* and *virtue* both can be understood as having a *character*-meaning (whereby they have some sort of standard *rule* that allows them to be respectively applied in different ways in different micro-contexts), hence, being cashed out with a specific *content*-meaning that is relative (i.e., situational) to those particular micro-contexts. If this is correct, then the generality problem and the situationist critique are not only misguided, but based on a misunderstanding of our actual epistemic practices.

BIBLIOGRAPHY

Adler, Jonathan & Michael Levin (2002). "Is the Generality Problem Too General?" *Philosophy and Phenomenological Research* 65:1 (July): 87-97.

Alfano, Mark (2011a). "Explaining Away Intuitions about Traits: Why Virtue Ethics Seems Plausible (Even If It Isn't)." *Review of Philosophy and Psychology* 2(1): 121-136.

---------- (2012). "Expanding the Situationist Challenge to Responsibilist Virtue Epistemology." *Philosophical Quarterly* 62(247): 223-249.

---------- (2013a). *Character as Moral Fiction*. Cambridge University Press.

---------- (2013b). "Identifying and Defending the Hard Core of Virtue Ethics." *Journal of Philosophical Research* 38: 233-260

---------- (forth.). "Knowledge, Right Action, and Virtue." In Mark Alfano & Abrol Fairweather eds., *Epistemic Situationism*. Oxford University Press.

---------- (forth.). "Virtues, Intelligences, and Situations." *Ethical Theory and Moral Practice*.

---------- (forth.). "What are the Bearers of Virtues?" In Hagop Harkissian & Jennifer Wright eds., *Advances in Moral Psychology*. Continuum.

Alfano, Mark and Abrol Fairweather (2014). *Epistemic Situationism*. Oxford University Press.

Alston, William (1995). 'How to Think About Reliability." *Philosophical Topics* 23:1 (Spring): 1-29

---------- (2005). *Beyond Justification*. Ithaca: Cornell University Press.

Alzola, Miguel (2008). "Character and Environment: The Status of Virtues in Organizations." *Journal of Business Ethics* 78(3): 343-357.

Annas, Julia (2003). "Virtue Ethics and Social Psychology." *A Priori* 2: 20-34.

---------- (2005). "Comments on John Doris's Lack of Character." *Philosophy and Phenomenological Research* 71(3): 636-642.

Annis, David (1978). "A Contextualist Theory of Epistemic Justification." *American Philosophical Quarterly* 15:3 (July): 213-219

Aristotle (1952). *Nichomachean Ethics*. Trans. by W.D. Ross. Great Books Series, Ed. Mortimer J. Adler. University of Chicago Press.

Axtell, Guy (2010). "Agency Ascriptions in Ethics and Epistemology: Or, Navigating Intersections, Narrow and Broad." *Metaphilosophy* 41(1): 73-94.

---------- (2014). "Thinking Twice about Virtue and Vice: On Character, Situations, and Dual Process Theory." In Mark Alfano & Abrol Fairweather eds., *Epistemic Situationism*. Oxford University Press.

Baehr, Jason (2006a). "Character in Epistemology." *Philosophical Studies* 128: 479-514.

---------- (2006b). "Character, Reliability, and Virtue Epistemology." *The Philosophical Quarterly* 56: 193-212.

Battaly, Heather (2008). "Virtue Epistemology." *Philosophy Compass* 3: 639-663.

Becker, Kelly (2008). "Epistemic Luck and the Generality Problem." *Philosophical Studies* 139: 353-366.

Beebe, James R. (2004). "The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis." *Noûs* 38: 177-195.

Bishop, Michael A. (2010). "Why the Generality Problem is Everybody's Problem." *Philosophical Studies* 151: 285-298.

Brandom, Robert (1994). *Making it Explicit*. Cambridge, MA: Harvard University Press.

---------- (2000). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.

Code, Lorraine (1987). *Epistemic Responsibility*. Hanover: Brown Univ. Press.

Cohen, Stuart (1987). "Knowledge, Context, and Social Standards." *Synthese* 73:1 (Oct.): 3-26

---------- (1988). "How to be a Fallibilist." *Philosophical Perspectives* 2: 91-123

Comesaña, Juan (2006). "A Well-Founded Solution to the Generality Problem." *Philosophical Studies* 129: 27-47.

Conee, Earl & Richard Feldman (1998). "The Generality Problem for Reliabilism." *Philosophical Studies* 89: 1-29

Darley, J.M. and C.D. Batson (1973). "From Jerusalem to Jericho: A Study of Situational and Dispositional Variables in Helping Behavior." *Journal of Personality and Social Psychology* 27: 100-108.

Davis, Wayne (2004). "Are Knowledge Claims Indexical?" *Erkenntnis* 64 (Nov.): 257-281

DeRose, Keith (1992). "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52:4(Apr.): 913-929

Doris, John M. (2002). *Lack of Character: Personality and Moral Behavior*. Cambridge University Press.

---------- (1998). "Persons, Situations, and Virtue Ethics." *Noûs* 32(4): 504-530.

Dretske, Fred (1981). "The Pragmatic Dimensions of Knowledge." *Philosophical Studies* 40:3 (Nov.): 363-378

Duncker, K. (1945) "On Problem Solving" *Psychological Monographs* 58 (American Psychological Association).

Estrada, C., A. Isen, and M. Young (1994). "Positive Affect Influences Creative Problem Solving and Reported Source of Practice Satisfaction in Physicians." *Motivation and Emotion* 18: 285-299.

Fairweather, Abrol & Owen Flanagan (2014). *Naturalizing Epistemic Virtue*. Cambridge University Press.

Feldman, Richard (1985). "Reliability and Justification." *The Monist* 68:2 (Apr.): 159-174.

Flanagan, O. (1991) *Varieties of Moral Personality*. Cambridge, MA: Harvard Press.

Funder, David C. (2006). "Towards a Resolution of the Personality Triad: Persons, Situations, and Behaviors." *Journal of Research in Personality* 40: 21-34.

Gettier, Edmund (1963). "Is Justified True Belief Knowledge? *Analysis* 23:6 (June): 121-123.

Goldman, Alvin (1976). "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73(20): 771-791.

---------- (1986). *Epistemology and Cognition*. Harvard University Press.

---------- (1992). *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Massachusetts Institute of Technology Press.

---------- (1993). "Ethics and Cognitive Science." *Ethics* 103:2 (Jan.): 337-360.

Gómez-Lobo, Alfonso. (1995) "Aristotle's Right Reason." In *Aristotle, Virtue and the Mean*, ed. R.Bosley, R.Shiner, & J.Sisson, pp.15-34.

Greco, John (2000). *Putting Skeptics in their Place: The Nature of Skeptical Arguments and their Role in Philosophical Inquiry*. Cambridge University Press.

---------- (2008). "What's Wrong with Contextualism?" *The Philosophical Quarterly* 58:232 (July): 416-436

---------- (2010). *Achieving Knowledge: A Virtue-Theoretic Account*. Cambridge University Press.

Harman, Gilbert (1999). "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error." *Proceedings of the Aristotelian Society* 99: 315-331.

---------- (2000). "The Nonexistence of Character Traits." *Proceedings of the Aristotelian Society* 100(2): 223-226.

---------- (2009). "Skepticism about Character Traits." *Journal of Ethics* 13(2/3): 235-242.

Heller, Mark (1995). "The Simple Solution to the Problem of Generality." *Noûs* 29:4 (Dec.): 501-515.

---------- (1999). "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives* 13: 115-129

Henderson, David (2009). "Motivated Contextualism." *Philosophical Studies* 142: 119-131

---------- (2011) "Gate-Keeping Contextualism." *Episteme* 8: 83-98

Hudson, Robert G. (2004). "The Generality Problem." *The Southern Journal of Philosophy* 42: 193-211.

Hurka, Thomas (2006). "Virtuous Act, Virtuous Dispositions." *Analysis* 66(289): 69-76.

Irwin, T.H. (1988). [Review of *The Fragility of Goodness*, by Martha Nussbaum]. *The Journal of Philosophy* 85:7 (July): 376-383.

Isen, A.M. and H. Levin (1972). "Effect of Feeling Good on Helping: Cookies and Kindness." *Journal of Personality and Social Psychology*, 21.

Isen, A., K. Daubman, and G. Nowicki. (1987). "Positive Affect Facilitates Creative Problem Solving," *Journal of Personality and Social Psychology* 52: 1122-31.

Jost, John T. & Lawrence J. Jost (2009). "Virtue Ethics and the Social Psychology of Character: Philosophical Lessons from the Person-Situation Debate." *Journal of Research in Personality* 43: 253-254.

Kamtekar, Rachana (2004). "Situationism and Virtue Ethics on the Content of Our Character." *Ethics* 114:3 (Apr.): 458-491.

Kaplan, David (2008). "On the Logic of Demonstratives." In *The Philosophy of Language*, 5th edition, A.P. Martinich, ed. New York: Oxford University Press.

King, Nathan L. (2014). “Responsibilist Virtue Epistemology: A Reply to the Situationist Challenge.” *The Philosophical Quarterly* 64:255 (Apr.): 243-253

Koehn, Glen (2008). “Character, Situation, and Intelligence.” *Proceedings of the XXII World Congress of Philosophy* 10: 245-260.

Kraut, Richard (2014). "Aristotle's Ethics", *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2014/entries/aristotle-ethics/>

Kristjánsson, Kristján (2008). “An Aristotelian Critique of Situationism.” *Philosophy* 83: 55-76.

Kupperman, Joel J. (2009). “Virtue in Virtue Ethics.” *The Journal of Ethics* 13: 243-255.

Lepock, Christopher (2009). “How to Make the Generality Problem Work for You.” *Acta Analytica* 24: 275-286.

---------- (2010). “Unifying the Intellectual Virtues.” *Philosophical and Phenomenological Research* 83(1): 106-128.

Lewis, David (1996). “Elusive Knowledge.” *Australasian Journal of Philosophy* 74(4): 549-567

Lukes, Stephen (2009). “Comment: Do People Have Character Traits?” In Chrysostomos Mantzavinos ed., *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice*, pp. 291-297. Cambridge University Press.

Machery, Edouard (2010). “The Bleak Implications of Moral Psychology.” *Neuroethics* 3:223-231.

MacIntyre, Alasdair (1984) *After Virtue*. Second edition. Notre Dame: University of Notre Dame Press.

Marmodoro, Anna (2011). “Moral Character versus Situations: An Aristotelian Contribution to the Debate.” *Journal of Ancient Philosophy* 5(2): 1-24

Mednick, M. (1963). “Research Creativity in Psychology Graduate Students.” *Journal of Consulting Psychology* 27: 265–6.

Merritt, Maria (2000). "Virtue Ethics and Situationist Personality Psychology." *Ethical Theory and Moral Practice* 3: 365-383.

---------- (2009). "Aristotelian Virtue and the Interpersonal Aspect of Ethical Character." *Journal of Moral Philosophy* 6: 23-49

Milgram, S. (1974). *Obedience to Authority*. New York: Harper and Row.

Miller, Christian (2003). "Social Psychology and Virtue Ethics." *The Journal of Ethics* 7(4): 365-392.

---------- (2009). "Empathy, Social Psychology, and Global Helping Traits." *Philosophical Studies* 142: 247-275.

---------- (2013a). *Character and Moral Psychology*. Oxford.

---------- (2013b). *Moral Character: An Empirical Theory*. Oxford.

Mischel, Walter and Yuichi Shoda (1995). "A Cognitive-Affective System Theory of Personality: Reconceptualizing Situations, Dispositions, Dynamics, and Invariance in Personality Structure. *Psychological Review* 102(2): 246-268.

---------- (1999) "Integrating Dispositions and Processing Dynamics within a Unified Theory of Personality: The Cognitive-Affective Personality System. In *Handbook of Personality: Theory and Research*, 2nd ed. Eds. L.A. Pervin and O.P. John. New York: Guilford, 197-218.

---------- (2008) *Introduction to Personality: Toward an Integrative Science of the Person*, 8th ed. Hoboken, NJ: Wiley.

Montmarquet, James (1987). "Epistemic Virtue." *Mind* 96: 482-497.

---------- (1993). *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman and Littlefield.

---------- (2003). "Moral Character and Social Science Research." *Philosophy* 78:3 (July): 355-368.

Morton, Adam. (2013). *Bounded Thinking: Intellectual Virtues for Limited Agents*. Oxford University Press.

Nussbaum, Martha (1986). *The Fragility of Goodness*. New York: Cambridge University Press.

Pollock, John & Joseph Cruz (1999). *Contemporary Theories of Knowledge*. 2nd Edition (1st ed., 1986). Lanham, Mass.: Roman and Littlefield Publishers.

Plantinga, Alvin (1993a). "Reliabilism." In *Warrant: The Current Debate*, pp. 182-210. Oxford University Press.

---------- (1993b). "Warrant: A First Approximation." In Warrant and Proper Function, pp. 3-20. Oxford University Press.

Putnam, Hilary (1983). *Reason, Truth, and History*. Cambridge Univ. Press.

Roberts, Robert & Jay Wood (2007). *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford University Press.

Rorty, Richard (1980). "Pragmatism, Relativism, and Irrationalism." *Proceedings and Addresses of the American Philosophical Association* 53:6 (Aug.): 719-738.

Ross, L. (1977). "The Intuitive Psychologist and his Shortcomings." In L. Berkowitz (ed.), *Advances in Experimental Social Psychology* v.10. New York: Academic Press.

Ross, L. & R. Nisbett (1991). *The Person and the Situation: Perspectives of Social Psychology*. New York: McGraw-Hill.

Russell, Daniel C. (2009). *Practical Intelligence and the Virtues*. Oxford Univ. Press.

Rysiew, Patrick (2001). "The Context-Sensitivity of Knowledge Attributions." *Nous* 35(4): 477-514.

Sabini, John & Maury Silver (2005). "Lack of Character? Situationism Critiqued." *Ethics* 115:3 (Apr.): 535-562.

Sarkissian, Hagop (2010). "Minor Tweaks, Major Payoffs: The Problems and Promise of Situationism in Moral Philosophy." *Philosophers' Imprint* 10:9 (Aug.): 1-15.

Simpson, Peter (1997). "Contemporary Virtue Ethics and Aristotle." In *Virtue Ethics: A Critical Reader*, ed. Daniel Statman, pp. 245-259. Washington, DC: Georgetown University Press.

Snow, Nancy (2009). *Virtue as Social Intelligence: An Empirically Grounded Theory*. New York: Routledge.

---------- and Franco Trivigno (2014). *The Philosophy and Psychology of Virtue: An Empirical Approach to Character and Happiness*. New York: Routledge.

Sosa, Ernest (1986). "On Knowledge and Context." *The Journal of Philosophy* 83(10): 584-585.

---------- (1988). "Knowledge in Context, Skepticism in Doubt: The Virtue of our Faculties." *Philosophical Perspectives* 2: 139-155.

---------- (1991). *Knowledge in Perspective*. Cambridge University Press.

---------- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge.* Volume I. Oxford: Clarendon Press.

---------- (2009). "Situations Against Virtues: The Situationist Attack on Virtue Theory." In Chrysostomos Mantzavinos ed., *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice*, pp. 274-290. Cambridge University Press.

---------- (2011). *Knowing Full Well.* Princeton University Press.

Sreenivasan, Gopal (2002). "Errors about errors: Virtue theory and trait attribution." *Mind* 111: 47-68.

---------- (2008). "Character and Consistency: Still More Errors." *Mind* 117(467): 603-612.

Swanton, Christine (2003). *Virtue Ethics: A Pluralist View*. Oxford.

Tweedale, Martin (2007). "Aristotle's Pragmatic Methodology for Ethical Enquiry." Unpublished.

Upton, Candace L. (2005). "A Contextual Account of Character Traits." *Philosophical Studies* 122(2): 133-151.

---------- (2009). "Virtue Ethics and Moral Psychology: The Situationism Debate." *Journal of Ethics* 13: 103-115.

Webber, Jonathan (2006). "Virtue, Character and Situation." *Journal of Moral Philosophy* 3(2): 193-213

---------- (2007a). "Character, Global and Local." *Utilitas* 19:4 (Dec.): 430-434.

---------- (2007b). "Character, Common Sense, and Expertise." *Ethical Theory and Moral Practice* 10: 89-104.

Wright, Sarah (2010). "Virtues, Social Roles, and Contextualism." *Metaphilosophy* 41: 95-114

---------- (2011). "Knowledge and Social Roles: A Virtue Approach." *Episteme* 8: 99-111.

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

---------- (2010). "Exemplarist Virtue Theory." *Metaphilosophy* 41 (1-2.): 41-57.

Zeller, Deborah (2007). "Virtue, Virtue Skepticism, and the Milgram Studies." *Philosophy in the Contemporary World* 14(2): 50-59.