

University of Alberta

Epistemic Value Pluralism

by

Michael Lockhart

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

Doctor of Philosophy

Department of Philosophy

©Michael Lockhart

Spring 2012

Edmonton, Alberta

Permission is hereby granted to the University of Alberta Libraries to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only. Where the thesis is converted to, or otherwise made available in digital form, the University of Alberta will advise potential users of the thesis of these terms.

The author reserves all other publication and other rights in association with the copyright in the thesis and, except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatsoever without the author's prior written permission.



Library and Archives
Canada

Published Heritage
Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et
Archives Canada

Direction du
Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-87890-3

Our file Notre référence

ISBN: 978-0-494-87890-3

NOTICE:

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

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

AVIS:

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

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

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

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

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

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

Canada

Dedication

To my wife Lori for all her love and support.

Abstract

Philosophers have, for the most part, taken for granted that all epistemic value is derived from the truth goal. Despite the recent development of virtue epistemology and its promise to reframe traditional problems, epistemic value monism remains largely unchallenged. I argue that once one conceptually prioritises agents over beliefs as virtue theories purport to do, value pluralism is implied.

In fact, monism becomes increasingly implausible once we acknowledge that the primary object of evaluation is a *situated, embodied, embedded, and bounded* agent.

My arguments take two forms. First, I distinguish between veridical virtues and virtues of epistemic usefulness. Where the former aim at the truth goal, the latter does not. I argue that there are both commonsense and paradoxical virtues of epistemic usefulness that intellectual exemplars exhibit that cannot be accounted for under monism.

Second, I argue that two prominent views in epistemology that claim to reject monism are suspiciously myopic and that a full commitment to pluralism is better-equipped to handle many traditional topics in epistemology such as the value problem, understanding, epistemic autonomy and responsibility, and wisdom.

Acknowledgement

I wish to thank my supervisor Adam Morton for the countless conversations, comments, criticisms, advice, wisdom, patience, and multitude of intellectual virtues from which I have learned.

Table of Contents

Introduction	1
Eudaimonia	4
Overview	5
Chapter 1	7
The Truth Goal: Two Prominent Theories	7
What is a Virtue?	20
Two Kinds of Epistemic Virtues	34
The Telos of Virtuous Epistemic Agents	38
Chapter 2	52
Commonsense	52
Knowing When One Has Enough Evidence	55
Knowing When One is in Over One's Head	72
Knowing When to Look for Empirically Adequate Beliefs	75
Knowing When to Forget	81
Knowing When to Ignore Relevant Low-level Problems	84
Scepticism	85
Excellence and Fitness	89
Chapter 3	92
Finite Agency and Long-term Success	92

Coupling Constraints	94
The Virtue of Dogmatism	96
The Virtue of Self-Deception	102
The Virtue of Intellectual Sloppiness	109
The Virtue of Dishonesty	113
The Paradoxical Nature	118
Chapter 4	124
The Value Problems	124
Understanding	144
Autonomy and Responsibility	151
Stupidity	156
Conceptual Priority and Community	158
Conclusion	162
Bibliography	168

Introduction

Since Ernest Sosa re-introduced intellectual virtues in his famous paper "The Raft and the Pyramid," epistemologists have employed them with substantial variety. Despite the diversity of uses, there are two general motivations for doing epistemology with virtues. The first is that many philosophers think the concept of virtues provides a powerful tool for handling traditional problems in epistemology. Sosa, for example, appealed to epistemic virtue as a way to overcome the debate between Foundationalism and Coherentism. Others find virtues helpful in addressing the Gettier problem, the lottery paradox, and resolving the debate between internalism and externalism. According to some epistemologists, virtues have a promising role in developing theories of justification, reliability, knowledge, and more generally, rationality.

A different motivation comes from a hesitation or dissatisfaction with the parochial nature of traditional epistemology. Instead of being confined to developing and refining theories of knowledge and justification, some employ virtue language to reframe the entire discipline by conceptually prioritising the agent before her beliefs. This widens the scope of epistemology by not assuming at the outset that true belief and knowledge are the central concerns of epistemology.¹ Instead, it is an open question as to what sorts of epistemic value there are, specifically whether there is more value than acquiring true beliefs

¹ It is, however, compatible with such a view.

and avoiding false ones. Wayne Riggs has called this approach *value-driven epistemology*.² The first approach has as its central question, "What is knowledge"?³ A value-driven epistemology seeks to answer broader questions such as, "What is epistemically valuable?" It may turn out that acquiring true beliefs and avoiding false ones are the only values, but a value-driven epistemology does not start with this assumption.

My approach is the latter one: value-driven virtue epistemology. In doing so, I take up what can be called an agent-based virtue epistemology. That is, agents take conceptual priority over beliefs in epistemic analyses and evaluations. Thus, I take one of the central tasks of virtue epistemology to be describing the kinds of characters, actions, skills, traits, properties, and processes involved in epistemic exemplars. Consequently, I wish to identify epistemic virtues in their own right, which may or may not be derived from their relation to truth. There has been some discussion about whether virtue epistemology makes this shift in priority by definition. Linda Zagzebski has argued that some views, most notably the varieties of virtue reliabilism put forth by Ernest Sosa and John Greco, claim to be virtue theories and yet place a conceptual priority on processes over the agent. Greco has responded to this accusation, insisting that priority is indeed placed on the agent.⁴ Regardless of what one makes of virtue

² Riggs, Wayne. 2008.

³ As well as questions such as, "What makes a belief justified, reliable, or warranted?"

⁴ Greco, John. 2000.

reliabilism, my own approach places conceptual priority on agents over other objects of evaluation such as beliefs.

Objects of evaluation are motivated by one's interests. For example, sometimes we care only to know if a belief is true. We may not be concerned with who gets credit or whether the belief was obtained from a reliable character. For my interests, agents take conceptual priority.

This approach allows knowledge, truth, reliability, and justification to take central roles in epistemology, but it also leaves it as an open question what sorts of epistemic value exist. Indeed it is my aim to argue that there are more epistemic values than acquiring true beliefs and avoiding false ones. By adopting an agent-based epistemology we can identify additional values that are properly epistemic. That is, some epistemic value is not derivative of truth. I call this view epistemic value pluralism; unless otherwise noted, hereafter called pluralism.

Why not argue for pluralism from a more traditional approach? First, as I have already suggested, virtue epistemology is conducive to identifying a variety of values. We have vast amounts of data because epistemic evaluations are already commonplace among both philosophers and non-philosophers. Ascriptions of epistemic virtues—Frank is open-minded, Sally is careful—are just as common as moral ones—Mary is courageous. And just as ethicists can identify moral values by studying the diversity of virtues found in exemplars, so too can epistemologists. Second, traditional approaches take for granted that all value is derived from truth, which is reinforced by conceptually prioritising beliefs. I

should note here that virtue epistemologists do not typically endorse value pluralism, even though some sound as if they do.⁵ Regardless, an agent-based virtue approach has a strategic advantage in that it does not dismiss pluralism out of hand.

Eudaimonia

Virtue epistemology has its roots, not surprisingly, in virtue ethics, both Aristotelian and Stoic. Some epistemologists, such as Linda Zagzebski, have been heavily influenced by Aristotle. Others are less influenced. For example, Sosa redescribes virtues more generally as capacities. However, almost all epistemologists define virtues by their relation to truth. I will argue that this is a mistake. In doing so, I will borrow from the ethicist's playbook. Just as philosophers have recognised the value of employing virtue theory to reorient moral theory, overcoming the restrictive conceptual framework of notions such as "obligation" and "duty", so too can virtue language allow us to reframe epistemology without being tethered to monism.

To that end, Aristotle's notion of eudaimonia is helpful. Eudaimonia has been understood by philosophers in a variety of ways. It is most commonly understood as flourishing, happiness, or long-term health and prosperity. This diversity of understanding is part of what I find attractive about the concept of

⁵ More on that later.

eudaimonia; it is open—perhaps even vague—yet informative. One can avoid restricting epistemology by asking what it means to flourish epistemically (as opposed to moral flourishing). It is possible that the answer is reducible to acquiring a high ratio of true beliefs. I will argue that it is not. The diversity and richness of our epistemic evaluations imply pluralism.

Overview

My task then is to argue that there are more epistemic values than traditionally identified. My arguments can be divided into two kinds. Chapters one and four argue more formally and abstractly for pluralism. Chapters two and three identify a variety of virtues used in epistemic evaluations that cannot be accounted for under monism.

In chapter one I discuss the nature of virtues, particularly epistemic ones. I briefly examine two theories—Ernest Sosa’s virtue reliabilism and Linda Zagzebski’s virtue responsibilism—and argue that they are ill-equipped to account for our epistemic evaluations. I then suggest an approach to virtue theory that allows us to overcome these shortcomings. In doing so, I distinguish between two types of epistemic virtues—the traditional truth-seeking kind and

virtues of epistemic usefulness.⁶ I conclude the chapter by proposing that the telos of agents is not reducible to acquiring true beliefs and avoiding error.

In chapter two I examine several commonsense virtues that imply value not derived from true belief or the avoidance of error. Some of these virtues are so commonplace we might describe individuals without them as having serious epistemic deficiencies or even psychological disorders. Yet these virtues cannot be accounted for under monism. I conclude the first chapter with some brief remarks about scepticism in light of commonsense virtues.

Chapter three continues to identify virtues that entail pluralism. Unlike the commonsense variety, here I focus on what Adam Morton has called “paradoxical virtues”. These virtues are unique in that they *sound like vices*.

In chapter four I argue for the inadequacy of monism to account for the value problem, understanding, epistemic autonomy and responsibility, and stupidity.

I conclude with some remarks on the necessity of adopting pluralism to account for wisdom as a distinctly epistemic aim.

⁶ These virtues will become clearer later on. For now I want to clarify that their value is not derived from pragmatic value alone, though they often have some practical aspect to them.

Chapter 1

The Truth Goal: Two Prominent Theories

There are at least two reasons epistemologists vary in their accounts of virtues. First, the concept itself is underdetermined. Second, epistemologists have emphasised different aspects of the agent's functioning. Despite the diversity of virtue theories, there appears to be widespread agreement that the goal of the virtuous epistemic agent is to acquire true beliefs and avoid false ones. Some philosophers treat these as two separate aims; others reduce it to just one goal. For example, Alvin Goldman claims that truth seeking and error avoidance can be reduced to one value: veristic value.⁷ Likewise, Marian David reduces the twin desiderata to a singular "truth goal".⁸ Thus, many epistemologists have endorsed value monism.⁹

Aristotle:

*(When) . . . thought is concerned with study, not with action or production, its good or bad states consist in being true or false. For truth is the function of whatever thinks . . .*¹⁰

*Hence the function of each of the understanding parts is truth; and so the virtue of each part will be the state that makes that part grasp the truth most of all.*¹¹

⁷ Goldman, Alvin. 2001.

⁸ David, Marian. 2001. 152.

⁹ This may be described as value dualism if one prefers.

¹⁰ Aristotle. 1139a27-30.

¹¹ Aristotle. 1139b11-13.

René Descartes:

*So today I have learned not only what precautions to take to avoid ever going wrong, but also what to do to arrive at the truth.*¹²

Roderick Chisholm:

*We may assume that every person is subject to a purely intellectual requirement—that of trying his best to bring it about that, for every proposition *h* that he considers, he accepts *h* if and only if *h* is true. One might say that this is the person's responsibility or duty qua intellectual being.*¹³

William Alston:

*Epistemic evaluation is undertaken from what we might call "the epistemic point of view." That point of view is defined by the aim at maximizing truth and minimizing falsity in a large body of beliefs . . . Any concept of epistemic justification is a concept of some condition that is desirable or commendable from the standpoint of the aim at maximizing truth and minimizing falsity . . .*¹⁴

Laurence Bonjour:

*What makes us cognitive beings at all is our capacity for belief, and the goal of our distinctively cognitive endeavors is truth . . . It follows that one's cognitive endeavors are justified only if and to the extent that they are aimed at this goal, which means very roughly that one accepts all and only those beliefs that one has good reason to think are true.*¹⁵

Jonathan Kvanvig:

The goal of the game is to find truth . . . Any claim that there are properties of belief that have value intrinsically, independent of any relationship to the truth, should be met with incredulity. . . . any property

¹² Descartes, René. 1984. 43.

¹³ Chisholm, Roderick. 1977. 14.

¹⁴ Alston, William. 1985. 83-84.

¹⁵ Bonjour, Laurence. 1985. 7-8.

*of belief that is valuable from a purely intellectual point of view had better find some connection between that property and truth.*¹⁶

In Marian David's article "Truth as the Epistemic Goal" he cites several other philosophers who appear to endorse monism: Paul Moser, Richard Foley, Keith Lehrer, Alvin Goldman, Ernest Sosa and Alvin Plantinga.¹⁷ It is worth noting that many philosophers talk *as if* they are pluralists. Unfortunately, they often do not adequately consider the consequence of rejecting monism. Sosa, for example, denies being a monist. But given that he defines virtues by their relation to veristic value alone, it is difficult to understand what other kinds of epistemic value he can acknowledge that are not reducible to the one.¹⁸ When virtue epistemologists do identify a plurality of values, they usually assume these are reducible to the veristic one. Nevertheless, monism appears to dominate the field. Since this is what I will challenge, let us identify the truth goal as the following: *I believe <p> if and only if <p> is true.* This should be understood as a double-aspect goal of trying to obtain true beliefs and to avoid false ones. So, for example, one can achieve this goal in situations by withholding belief to avoid error.¹⁹

My goal in this chapter is to give a broad treatment of virtue theory that allows us to identify epistemic virtues commonly observed in agents, but not

¹⁶ Kvanvig, Jonathan. 2003. 54.

¹⁷ See David's endnote for specific references to those philosophers' works.

¹⁸ In chapter four I discuss this further.

¹⁹ Although one may have an obligation to seek more evidence in order to form some belief. See Hall, Richard J and Charles R. Johnson. 1998. 129-139.

accounted for by monistic theories. There are probably other useful ways of describing virtues. I am not claiming my approach is the best, only that it is more useful than some others for doing value-driven epistemology. Before briefly describing what I take to be a useful and attractive account of virtues, I will mention two prominent theories.

The first was put forth by Ernest Sosa. Contemporary virtue epistemology owes much of its existence and interest to Sosa. The second theory I will consider is from Linda Zagzebski, which has also been quite influential. Where Sosa's account puts reliability at its centre, Zagzebski's emphasises the responsibility of the agent. It is worth noting that nothing in their respective accounts limits the role of responsibility and reliability. In fact, Zagzebski insists her account places great importance on reliability.²⁰ And Sosa could just as easily build into his account a more prominent role for responsibility. Nonetheless, Sosa and Zagzebski are appropriate representatives of what have become known as virtue reliabilism and virtue responsibilism respectively. Let us look first at Sosa's account of virtues.

Sosa's theory continues to be developed in several of his writings since he first invoked epistemic virtues in "The Raft and the Pyramid". Sosa's definition of an epistemic virtue can be summarised as follows: *Virtues are dispositions to*

²⁰ Some have questioned this.

*have true beliefs about propositions in a field when in a particular environment under particular conditions.*²¹

Virtues are properties of agents. On Sosa's account, they are dispositions that arise from the subject's inner nature.²² It is not surprising then that *internal* justification plays an important role for Sosa, for example, in addressing both the new evil-demon problem and the problem of meta-incoherence.²³

However, Sosa is not simply an internalist. While virtues reside in the subject, they are defined relative to particular environments. An agent's environment is constituted by a complex set of conditions and properties, not simply time and location. Indexing virtues to a particular environment *E* means that for a subject *S* to possess a particular virtue *V*, *S* does not need to be in *E*. *S* can be in *E'* and still retain *V*.²⁴

For example, suppose Mary is intellectually virtuous; she has admirable capacities of perception, memory, reasoning, and problem solving. She cautiously gathers evidence, resists hasty conclusions, forms beliefs based on what she takes to be good empirical evidence, and has confidence in her own cognitive abilities. Because these are epistemically virtuous, she is disposed to acquire more true beliefs than false ones in an everyday natural environment.

Now suppose Mary is placed in a classic demon environment of illusions. What her cognitive faculties tell her are real objects are in fact not. Using the

²¹ Sosa, Ernest. 1991. 140.

²² Sosa, Ernest. 1991. 141.

²³ More on that below.

²⁴ Sosa, Ernest. 1991. 140.

same cognitive abilities listed above, Mary acquires more false beliefs than true ones. But Sosa insists we should not withhold ascribing to Mary the virtues we had prior to her insertion into the demonic environment. Her powers of reasoning, perception, memory, and problem solving are still admirable.²⁵ She remains an intellectually virtuous agent. Why? While virtues are contextual, one does not need to be in the relevant environment in order to have the virtue. A subject, *S*, has *V* even if in *E'* because *S* meets the conditional that if *S* were in *E*, *S* would have a high ratio of true beliefs. So even though Mary acquires many false beliefs in her demon environment *E'*, she is still virtuous because she meets the conditional that if she were in a more natural environment *E* she would have acquired a high ratio of true beliefs.²⁶ One might attempt a lengthy description of what counts as a virtue-endorsing environment. But since there is a wide variety of virtues, giving a general description of appropriate environments will prove difficult. For my purposes, suffice it to say that Sosa has little interest in bizarre environments such as the demon one. His interests lie in more “naturally” occurring ones.²⁷

Notice that so far on Sosa’s account one can acquire false beliefs virtuously. This is for two reasons. First, one may be exercising the wrong virtue

²⁵ That is not to say she should use those virtues *in that environment*.

²⁶ On Sosa’s view, Mary would also enjoy what he calls internal justification. Sosa’s own account of justification is also indexed, though more generally he maintains the indexing is to our actual natural environment *E*, not the demon’s environment *E'*. This indexed notion of justification is central to his reply of the new evil demon problem. Mary enjoys intellectual virtue as well as internal justification because both of those notions for Sosa are indexed to environments that are not *E'*. See Sosa, Ernest. 1991. 143-144.

²⁷ What counts as a naturally occurring environment is something the sceptic might challenge.

for the occasion. Second, virtues are not guarantors of true beliefs. One may, for example, have a virtue in the correct environment but simultaneously have a vice that prevents one from acquiring the truth. Or one might be unlucky. Virtues are only dispositions toward acquiring true beliefs. This can be understood in terms of reliability. "X is an intellectual virtue only if x would produce a high ratio of true beliefs."²⁸ And of course, Sosa intends this high ratio to be indexed to a particular environment. Having now explained disposition in terms of reliability, we can see why Sosa's view is a form of virtue reliabilism.²⁹

Before moving on to Zagzebski's account, I wish to make one final comment regarding environments on Sosa's account. Sosa handles the new evil-demon problem above by indexing both virtues and justification to *our* natural world. However, one should not infer from this that on Sosa's account all virtues must be indexed to "normal" environments. There are at least two reasons this interpretation is unfavorable. First, one will be hard pressed to find a satisfying description of "normal" environments and sets of conditions that are not also so vague as to be uninformative. Our everyday environments are too varied; and unique conditions are perfectly natural. Second, and more importantly, some virtues, by their very nature, cannot be indexed to "normal" conditions. Some must be indexed to extraordinary or irregular environments and peculiar conditions.

²⁸ Sosa, Ernest and Laurence Bonjour. 2003. 156.

²⁹ Other virtue reliabilists include John Greco and Alvin Goldman.

Take for example the capacity to navigate one's environment during times of low or nil visibility. This may be during snow storm conditions similar to arctic whiteouts or a power outage in one's home. Many intellectual capacities are manifest only in these rare—yet natural—conditions; and the heuristics used in such cases would be less than optimal under so-called "normal" conditions.

Other virtues depend on shifting environments and conditions. So environments should not be understood as static ones. Take for example the virtue of being flexible. Yang may be particularly skilled at working a room, regardless of the room he is in. He may exercise this skill in Toronto one day, Tokyo the next, and then still later in Tijuana. He knows when to speak, to listen, to be aggressive, to show humility, and what to pick out as salient—culturally relevant—information. These environments, while all appropriately described as natural and social, are rather diverse. They vary in degree of social norms. And the virtue being picked out is one that is dependent on shifting environments: Yang's flexibility to network in a variety of social settings.³⁰ So indexing to environments and conditions ought to be understood to include irregular and dynamic ones. We could also think of virtues that might arise in extremely rare, or even impossible, environments. For my purposes, I will set these aside. The virtues I am interested in are ones we find in real environments, not, for example, demon environments.

³⁰ While clearly a social virtue, Yang is sifting through, interpreting, and manipulating a substantial amount of information coming to him. Surely this involves epistemic capacities.

While Sosa places great emphasis on reliability, others such as Zagzebski emphasise the agent's responsibility. In her influential work *Virtues of the Mind*, Zagzebski attempts to develop an Aristotelian virtue epistemology. She argues that epistemic virtues can be modeled, more or less, on those found in ethics. In fact, her project is rather ambitious in that she seeks a unity of these virtues, arguing that epistemic virtues are closely related to moral ones. It is no surprise then that she puts forth a definition of virtues that is supposed to hold in both domains.³¹ A virtue is: *A deep and enduring acquired excellence of a person, involving a characteristic motivation to produce a certain desired end and reliable success in bringing about that end.*³²

For my present purposes, I am interested in how she applies this definition specifically in epistemology. There are several things to note about Zagzebski's concept of epistemic virtues. First, she says that virtues are deep and enduring acquired excellences of a person. By "deep and enduring" Zagzebski means not only that virtues are properties of agents, but that they are so intimate as to be closely associated to one's identity.³³ This differs from Sosa's account in which virtues are also properties of the agent but do not necessarily constitute the agent's personality or identity. They only establish the agent's reliability.³⁴ While there is some debate on how closely we ought to relate virtues with identity, we can observe that as a matter of practice we often do

³¹ She also claims this is a definition that applies to other domains such as the religious one.

³² Zagzebski, Linda. 1996. 137.

³³ Zagzebski, Linda. 1996. 85, 125- 126, and 135. Zagzebski maintains virtues are states of the soul.

³⁴ John Greco seems to agree with Sosa on this issue.

associate personalities with virtues. Sometimes an attribution of virtue is tantamount to a description of that agent's personality. Take the following common ascriptions:

Jones is open-minded.
Jo is imaginative.
Arthur is careful.
Marie is witty.
John is generous.
Stephanie is courageous.

These ascriptions pick out more than temporary properties; they pick out personalities or character traits.³⁵ For example, we do not usually consider someone open-minded if she displays an open mind on one occasion while being dogmatic or stubborn the majority of the time. Likewise, we would not say Stephanie is courageous if she is typically a coward, even if we acknowledge that some token act took courage to perform.

That is not to say that people lacking a particular virtue cannot display it occasionally. Stephanie may be a coward most of the time and still display courage in some particular circumstance. This is presumably how people acquire virtues. People are not born with all of the virtues. Following Aristotle, Zagzebski maintains virtues are acquired through repetition and habituation. It is "part of the nature of a virtue...that it be acquired by a process of habituation."³⁶

³⁵ Sometimes a person's personality is cashed out in terms of vices, for example, "Mary is greedy".

³⁶ Zagzebski, Linda. 1996. 125

Furthermore, on Zagzebski's account, what is acquired is an excellence; virtues are thought to be good or valuable.³⁷ Thus, *virtue* is a term of praise and *vice* is one of blame or shortcoming.

Contrary to a teleological approach, Zagzebski maintains a motivation-based theory. Virtues involve a "characteristic motivation...to produce a certain desired end." A teleological approach holds that virtues are aimed at a particular end. On Zagzebski's account, a motive is not simply a desired end; it is also an emotional state, even if weak in its intensity. For example, the virtue of benevolence involves "the disposition to have characteristic emotions that direct action in a particular direction, probably the well-being of others."³⁸ The motive will be different for different virtues. So the virtue of being fair-minded will have one distinct motive; courage will have a different one. On this account, a "motive is an action-directing emotion" and is always present when virtues are being employed.³⁹

With regard to intellectual virtues in particular, Zagzebski maintains, "intellectual virtues can be defined in terms of motivations arising from the general motivation for knowledge..."⁴⁰ The motivation for knowledge is,

³⁷ It is worth exploring the sense in which virtues are supposed to be good. After all, it is not clear they are good in every sense. For example, they are not always beneficial for the agent to have. It may be that virtues are good not because they benefit the agent, but because they are admirable. Zagzebski discusses this in *Virtues of Mind*.

³⁸ Zagzebski, Linda. 1996. 132.

³⁹ Zagzebski, Linda. 1996. 130-132.

⁴⁰ Zagzebski, Linda. 1996. 166.

according to Zagzebski, intrinsically good.⁴¹ It does not derive its goodness from anything else, even the goodness of knowledge.⁴² In fact, the goodness of intellectual virtues⁴³ is derived from the agent's motivation for knowledge.

Finally, Zagzebski maintains that virtues are reliably successful at bringing about appropriate ends. In the case of intellectual virtues, that end is knowledge.

It is worth noting that, for Zagzebski, the motivational and reliability components are not simply independent conditions for intellectual virtues. Since the motive for knowledge is intrinsically good, and the motivational component of each intellectual virtue is derived from this motive, she finds it reasonable that "reliable success in achieving the aim of a good motive is itself a good thing."⁴⁴ It then follows that the goodness of the reliability component of an intellectual virtue is derived from the goodness of the motivational component. So both conditions are related through goodness; and they are both agent-based.⁴⁵

Zagzebski's account appears more demanding on the agent overall compared to Sosa's account. For example, Zagzebski requires that the agent has a motivation associated with the virtue. Sosa's account includes no such requirement. According to Zagzebski, there is a close relation between virtues

⁴¹ Zagzebski, Linda. 1996. 209.

⁴² Zagzebski, Linda. 1996. 202-203. This is consistent with her virtue ethics in that the goodness of a motivation, for example the motive to bring about the well-being of others, is not derived from anything else that is good, including the well-being of others.

⁴³ I think a case can be made to distinguish intellectual virtues from epistemic ones. Perhaps the latter is employed in belief-forming contexts and the former are used in intention-forming ones. Since the literature does not consistently make the distinction, neither will I make it in the present work.

⁴⁴ Zagzebski, Linda. 1996. 209.

⁴⁵ Zagzebski, Linda. 1996. 209.

and personal identity. By contrast, Sosa maintains that virtues might be nothing more than a reliable capacity or faculty of the agent. While reliability plays a role in both accounts of virtue, it differs in a rather substantial way. Zagzebski's reliability component is related to the motivational one. Reliable success on her account is aimed at *knowledge*, since that is what the agent is motivated towards. However, Sosa thinks that virtues are reliable, when properly indexed, in acquiring *true beliefs*. Despite this difference in definition, both philosophers maintain that knowledge requires the use of virtues.

Both philosophers have been immensely influential. They have done much to set the landscape of virtue epistemology. And while they differ on some rather substantial points, they both define virtues in terms of reliability and truth conduciveness.⁴⁶ They are in good company. Reliability in obtaining true beliefs is a feature of many virtue theories.⁴⁷

Unfortunately, theories that define virtues in terms of truth conduciveness fail to acknowledge some virtues we see manifest regularly by epistemic exemplars, since some virtues are not aimed at acquiring true beliefs. Sometimes they are aimed at false ones. Others may be aimed only at empirically adequate or useful beliefs. Such virtues will be discussed later. For now, we can note that monistic accounts must classify any virtue that does not aim at the truth goal as non-epistemic. My task is to argue that some virtues are

⁴⁶ For Zagzebski, the aim is knowledge.

⁴⁷ Not all virtue theorists give such a central role for reliability. James Montmarquet, for example, rejects reliability as a necessary condition of virtue.

epistemic even though they do not aim at veristic value. If I am correct, no theory ought to include truth conduciveness as a necessary component of epistemic virtues.

I employ two strategies. The first is more abstract and formal. I argue that not all epistemic value can be reduced to acquiring truth beliefs and avoiding false ones. The second strategy identifies virtues that can be identified as properly epistemic only if we assume there is more value than the veristic one. Through these two approaches I urge us to widen the scope of epistemic value. I argue this follows from an agent-based epistemology. Once we consider the range of evaluations we make of epistemic exemplars,⁴⁸ it is necessary to acknowledge a plurality of value. Consequently virtue theories such as the two above are problematic since they define virtues by only one value. Instead, we need a theory that allows for a plurality of value. So before presenting my case for rejecting monism, I will sketch a theory that accommodates the variety of virtues we see epistemic exemplars displaying.

What is a Virtue?

Philosophers have varied in their accounts of virtues. However, what they all appear to have in common is that they think virtues are good things, either intrinsically or instrumentally. The goodness of virtues is where the universal

⁴⁸ We could also identify epistemic value by evaluating those who fail miserably. It is reasonable to think that we ascribe failure and vice based on more than the agent's inability to obtain true beliefs.

agreement appears to end. In fact, philosophers do not even agree on what it means for virtues to be good.

I want to emphasise one quick point before giving a brief description of features I suspect need to be identified of virtues to allow for pluralism. The term "virtue", like many ancient words, carries historical baggage and is used in substantially different ways today. As mentioned above, its various uses are due, at least partly, to different functions that has been emphasised. For example, we sometimes hear that "patience is a virtue", where virtue picks out a character trait or aspect of one's personality. Other times we see people using virtue language to denote general assets. It is not clear to me any one understanding or use of the term is more appropriate. For better or worse, the specific meaning of the word appears to be highly contextual. However, perhaps some uses of the term are better equipped to handle different domains. For example, treating virtues as character traits might serve in an account of morality, but not knowledge. Perhaps a broader sense of virtue is more appropriate for the latter. Maybe there is a unified understanding that best explains every domain and function, but I do not wish to argue for such a unity.

Consequently, when I examine Sosa's and Zagzebski's definitions of virtues, I am not suggesting their use of the term is altogether mistaken,⁴⁹ only

⁴⁹ It may be the correct understanding given what their theories are aimed at: a theory of knowledge.

that neither allow us to account for the richness of epistemic evaluations. Insofar as their interests lay elsewhere, their particular accounts may suit them just fine.

The two theories above give competing explanations of virtue. On the one hand, Sosa's account treats virtues as nothing more than faculties or capacities of the agent. Virtues are internal to the agent, but not necessarily something for which the agent is responsible.⁵⁰ For example, having an open mind and good eyesight are equally epistemic virtues on his account. Moreover, Sosa does not think virtues require a motivational component. What matters is that the virtue is internal or attributable to the agent and that it is reliable in a particular environment.

Things get more complex when Sosa applies his concept of virtue to his theory of knowledge. Sosa's own kind of virtue epistemology, Virtue Perspectivism, identifies two types of knowledge: animal and reflective. Animal knowledge is little more than the use of one's virtues in the appropriate context; he calls this apt belief. Reflective knowledge requires that one be aware of the aptness of one's belief; that is, it must be an aptly apt belief. We need not examine his theory of knowledge. I mention it here only to note that his account of virtues is motivated by and plays an important role in his theory of knowledge.

Virtues in Zagzebski's account are not simply faculties or assets, but instead traits of character. They are closely associated with one's personality and

⁵⁰ At least not in a robust sense. Sosa still maintains the agent deserves credit for well-functioning natural faculties like good eyesight even if she did nothing to have that asset.

identity. Virtues also require an appropriate motivational component as we saw above. And they are acquired over a long period of time through repetition and habit. So there will be a historical-causal account for any particular virtues. Finally, similar to Sosa, Zagzebski maintains a reliability component in her theory.

There are of course many other accounts of virtue that have been proposed in the literature. But again, these represent the two general approaches to virtues: reliable faculties and responsible traits of character.

As I develop the case for epistemic pluralism, it will be clear that neither of these views is able to account for the richness of our epistemic evaluations. However, I offer a few reasons below to doubt that the above accounts of virtues will suffice. Consider Zagzebski's account.

First, it is dubious that we need to acquire every virtue by habit or that an agent needs to have any particular causal-historical appropriation of her virtues. Maybe *some* virtues are acquired by habit and have a particular history associated with them. But neither feature is required to possess virtues. Take for example Swampman.⁵¹ Swampman pops into existence from some extraordinary events. By chance he has traits that in any ordinary human would be identified as open-minded, prudent, intellectually courageous, and charitable; and he displays impressive powers to efficiently gather evidence to form beliefs. Though he has no history and no habits, I see no reason we would withhold identifying these properties as virtues. And should he arrive at a true belief *because* he uses these

⁵¹ Davidson, Donald. 1987. 441-458.

virtues, then I think it is just as reasonable to ascribe knowledge to him.⁵² I see no reason to think humans ought to be evaluated any differently; some virtues are habitually acquired while others may be accidentally or abruptly learned. We might even be born with some virtues. John Greco points out⁵³ that this appears to be what Aquinas had in mind when he distinguished between habitual and non-habitual virtues,

Virtue designates a certain kind of perfection of a power. . . . Now some powers are of themselves determined to their acts, for instance, active natural powers. Hence these natural powers are in themselves called virtues. But the rational powers, which are proper to man, are not determined to some one thing, but are related indeterminately to many, and they are determined to their acts by habits, as we have said.⁵⁴

Aquinas seems to recognise that some virtues—rational powers—are habitually acquired while other virtues—natural powers—are not.

Two further things are worth noting about the above passage from Aquinas. First, it appears virtues were not always thought to be character traits as Zagzebski treats them, but rather sometimes identified with faculties, for example, sight, hearing, and memory. Defending her own account, Zagzebski points out that she is in line with the Aristotelian tradition that refrained from identifying natural capacities with virtues, preferring instead to view them as character traits. But some philosophers have pointed out that Aristotle was in the minority.⁵⁵ Many seemed to identify virtues with natural capacities, not only

⁵² See Sosa, Ernest. 2001. 56.

⁵³ Greco, John. 2000. 180.

⁵⁴ Aquinas, Thomas. *Summa Theologiae, 1-11, Question LV*. Translated in Oesterle, John. 1984.

⁵⁵ See Greco, John. 2000. See also Annas, Julia. 2003.

the perfection of those capacities. Others, like the Stoics, thought of virtues as skills. Second, a motivational component does not appear necessary for all virtues. The use of natural powers, which are identified by Aquinas as virtues, often occur without any motivation. Even if one rejects the faculty approach to virtues, some virtues seem to be passive traits. Some *happen to the agent*, rather than the agent making a conscious decision to employ them.⁵⁶ Passive traits do not require any robust sense of motivation. In the case of epistemic virtues, Zagzebski's view includes an even more stringent condition; they must be motivated toward *knowledge*. This requirement becomes even more dubious if one maintains, as Zagzebski does, that the use of virtues is necessary for knowledge. Perceptual knowledge, for example, does not require any corresponding motivation. So I find the motivational component rather problematic insofar as one wishes to employ it to define all virtues.

Some virtues *are* habitually acquired character traits that are motivated toward knowledge. However, I see no reason to think this is true of every virtue; and virtues have enjoyed a much broader identification both historically and in contemporary use. Given my current purposes, I favour a less restricting virtue theory.

Sosa puts much less demand on the notion of virtue, and so avoids the above criticisms. However, his concept of virtue is perhaps too thin. Recall that for Sosa, a virtue is nothing more than an asset possessed by the agent such that

⁵⁶ For example, being patient, intellectually charitable, and forgiving.

its use reliably produces true beliefs. But is that really all we mean when we attribute a virtue to someone? Are we not saying something more substantial in that we take virtue to be a term of praise? And yet we do not praise the agent just for having an asset like eyesight or hearing. Such a thin account captures only part of the content in virtue ascriptions.

There is disagreement in scholarship as to what the ancients had in mind when they referred to virtues. There is textual evidence to think that some ancients, for example Plato and then later Aquinas, identified virtues with natural capacities such as sight and hearing. John Greco argues that this is how we ought to understand these philosophers. However, Zagzebski argues against such exegetical views. She maintains that even non-Aristotelian accounts never intended the mere use of these capacities to be identified as virtues. When Aquinas wrote "Virtue designates a certain kind of *perfection of a power*"⁵⁷, he was not identifying virtues with natural capacities but rather the perfection of those capacities. Regardless of who is correct exegetically, the discussion⁵⁸ itself allows us to identify two essential properties of virtues: credit and praiseworthiness. But Sosa's account does not capture those properties.

People are not credited or praised simply for having good eyesight, being tall, or having green eyes. Yet it is clear other natural assets throughout history have been identified as virtues, for example, strength and beauty. Are these

⁵⁷ Emphasis mine.

⁵⁸ For more on this debate, see the exchange between John Greco and Linda Zagzebski in *Philosophy and Phenomenological Research* Vol. LX, No. 1, January 2000. I am sceptical we should be searching for an exegesis that treats this term in a univocal sense.

counter-examples helpful to Sosa? I do not think so. Some natural assets, but not all, deserve credit and praise. For example, the beauty we identify as virtuous requires a certain level of effort, insight, and self-grooming. Likewise, when we praise someone for their strength and fitness, we recognise that the agent has made choices leading to a healthy life; and that says quite a bit about the person living in a culture where obesity and related diseases are becoming more common. What distinguishes virtuous from non-virtuous natural assets is the role responsibility plays.⁵⁹ An agent is not responsible for her height, so it makes no sense to credit or praise her for being tall. But an agent *is* responsible for being healthy, strong, and well-groomed.⁶⁰ Likewise, when attributing a vice to someone, we are crediting and blaming an agent, even if minimally. And Sosa's account is problematic for this very reason. On his view, virtues include capacities for which we deserve no credit or praise. What sounds even more bizarre is that his theory treats deficient faculties, for which we deserve no blame, as vices.

The variety of epistemic evaluations we make requires a virtue theory that is not as stringent as Zagzebski's or as thin as Sosa's. Similar to Zagzebski and Sosa, I will sketch out an approach to virtue theory that follows ancient Greek thinkers, specifically the Stoics. According to the Stoics, virtues were skills

⁵⁹ This does not entail a component of motivation or intention; an agent can be responsible regardless of the presence or absence of motivation or intention. For example, an agent may be held responsible for perceiving something even in the absence of any motivation.

⁶⁰ Of course, some people have an advantage securing these natural assets. One might enjoy a natural beauty, good genes, a strong desire for physical activity and healthy foods. But these advantages are true of many virtues.

of living well.⁶¹ Concerning moral living, they were skills to achieve the good life. My approach also treats virtues as skills, which is helpful to avoid the shortcomings of the two views above.

I want to make a quick distinction between skill and capacity. I take capacity to mean more or less what Sosa does. It may be nothing more than a passive property, feature, or ability possessed by the agent. Capacity is more or less synonymous with asset. On the other hand, a skill is something that an agent deserves credit for possessing and using appropriately. For example, eyesight is a capacity. As long as an agent has her eyes open, she is going to see. That is just what it is to have the capacity of eyesight. However, being visually acute is a skill. Skills are not simply capacities, but rather, the ability to *use* those capacities in a praiseworthy manner.

This approach captures what Sosa's account does not: the credit—praise and blame—component of virtue and vice ascriptions. Moreover, a stoic approach can accommodate the attractive aspect of Zagzebski's view that virtues are closely associated with traits of characters without requiring some particular motivation, history, and acquisition through habit. A person's identity is often associated with her virtues, for example, "Mary is an open-minded *person*". Mary has a particular disposition toward her beliefs and cognitive faculties such that it makes up part of her personality. However, notice that on the stoic approach there is no need to insist Mary became open-minded through habit.

⁶¹ Annas, Julia. 2003. 26. Zagzebski follows Aristotle in rejecting the idea that virtues are skills.

And she might have no corresponding motivation when employing this virtue. Nevertheless, she may be an extremely open-minded person. The same can be said of patience. Perhaps Mary is naturally disposed to patience, a skill recognised, admired, and sought after by people who know her. While skills do not always pick out character traits, they can.⁶²

Interestingly, this approach is helpful to epistemologists who maintain that virtues are necessary for knowledge. If virtues are required for knowledge, they cannot be only traits of character, since one can gain knowledge even when acting *out of character*. Likewise, knowledge appears to be possible even if one has no motivation for it. For example, it is dubious how much character or motivation is present in simple perception, yet surely one gains knowledge sometimes through such processes.⁶³

While virtues in general can be thought of as skills of living, I wish to identify *epistemic* virtues with a particular kind of skill: skills of judgement. Specifically, they are skills in judgement about how to use one's intellectual capacities. Skills depend on several factors, including the environment in which an agent finds herself. As on Sosa's view, virtues on this account are indexed to environments and require the agent to have the skill in judgement about when and to what degree her capacities ought to be used. In environment *A* one may judge that *x* amount of caution is appropriate whereas in environment *B* *x* is

⁶² Creativity is also recognised as a skill acquired that often is related to one's identity.

⁶³ I am not endorsing the view that knowledge requires virtue. I only point out here that driving a wedge between virtues, traits of character, and motivation is useful for those that do maintain such a view.

inappropriate. The virtuous agent does not merely possess a cautious nature; she knows when to use the right amount of caution for the occasion.⁶⁴ While I will have more to say on this later, it is worth noting here that both the present account and the stoic one recognise the role practicality plays in virtue ascriptions. It should become clear that this follows from doing agent-based epistemology.

Consequently, any plausible account of epistemic virtues must include the following features. First, epistemic virtues are essentially *situated*. That is, they are skillful judgements that take place in the context of real-world environments, including real-time. This disqualifies, for example, demon environments such as the one mentioned above. Second, virtues are *embedded*. That is, virtues are constrained by the interaction between the physical agent and the world. Given that agents are embedded, they are limited in their possible capacities and behaviours. Third, virtues are *embodied*. The embodiment of the agent plays a pervasive role in the agent's outlook, interaction, and navigation of the world. Therefore any view of skills must also acknowledge her embodied nature. Finally, when indexing to environments one must include the *boundedness* of the agent. I will say more about these conditions below. For now, it suffices to say that agents have limited cognitive

⁶⁴ This is not unlike Aristotle's account of virtue as the mean between extremes.

capacities; so a virtue theory ought to reflect that fact.⁶⁵ The stoic approach accommodates these features.

There is one final aspect of the stoic view I wish to raise here. And it is very much related to what I have mentioned above concerning the human condition. In moral theory, virtue is taken to be a success term. So, for example, the morally virtuous person who tries but fails to stop the bank robber does not fail, morally speaking. She was courageous even though the thief got away. The success of the thief does not entail a failure of the agent in a distinctly *moral* sense. Even though the virtuous person will sometimes fail to stop the particular crime or tragedy, she may have achieved a moral accomplishment. So virtue theories are not typically consequentialist in nature; that is of course why some philosophers are attracted to virtue ethics in the first place.

However, if virtue is taken to be a success term, then what are we to make of Sosa's and Zagzebski's theories? On their accounts success is tied to the reliability of the agent to get the truth and avoid error.⁶⁶ But what of virtuous agents who fail in their immediate task? Think for example of a courageous person in a country of oppression. Suppose she fails to stop unjust behaviour. Nevertheless, surely she is courageous. That she was courageous entails she was successful. What about an agent who does not achieve the desired intellectual end of obtaining a true belief? On their accounts, even though the agent might

⁶⁵ I am aware that each of these features play prominent roles in cognitive science, decision theory, and feminist epistemology to name a few. It is not my intention to endorse any particular view in those areas here.

⁶⁶ See Greco, John. 2000.

display the virtue, she is not successful. But this evaluation betrays the success component of virtues. If an agent is acting virtuously, she is supposed to be succeeding. Virtue epistemologists may concede here that this is simply an asymmetry between moral and intellectual virtues. Where moral virtues are always successful, intellectual ones are not. Given the variety in which virtues have been used, this might not be too much cause for concern. However, there is no need to invoke such an asymmetry here between the two domains. There is a more natural way to understand both virtue and success.

According to Sosa and Zagzebski, virtues produce a particular consequence reliably: true belief (or knowledge). What this means is that an agent can be virtuous and still fail, which creates the asymmetry between ethics and epistemology. However, once we consider that our epistemic evaluations extend beyond beliefs, we can keep the symmetry: an agent acting virtuously in some environment x is always successful.

The symmetry is kept by acknowledging two types of aims that virtuous agents have: a token and a type aim. A token aim is an immediate or singular one whereas a type aim is some long-term goal. And while we sometimes evaluate someone with an eye on the token aim, we acknowledge that the type aim is often at least as equally important.⁶⁷ Consider first how this distinction has been employed in the moral domain. The virtuous agent who tried to stop the bank

⁶⁷ Sometimes we care more about the token aim. But even in such cases we still acknowledge type aims are significant.

robber may have failed in her token aim: stopping the robbery. However, she succeeded in her type aim of being a particular kind of person: a moral person or the kind of person that is not passive in the face of oppression. It is important to point out here that the type aim—cultivating moral character—is not essentially consequentialist. It does not depend on particular token outcomes.⁶⁸

Notice that both accounts above define virtues only by their token aims. While consistent with monism, it makes the use of virtue language rather awkward. Sometimes epistemic exemplars are identified, not by a single goal they achieved, but rather, because they achieved a long-term one; they became a particular type of inquirer. Both our current moral and epistemic evaluations already acknowledge these two tiers of success. The moral agent above was courageous. She exhibited a virtuous character and in doing so contributed to living the morally good life. Her token failure to bring about her desired state of affairs is hardly insignificant; but it is not the only thing we are evaluating; it is not even the primary thing.⁶⁹ The same is true in epistemology. A creative and honest scientist might endorse theories that available evidence supports but which are in fact false.

Once we acknowledge the importance of long-term success—becoming particular types of inquirers—it naturally follows that there is more to becoming virtuous epistemic agents than achieving token goals of acquiring true beliefs

⁶⁸ Admittedly, it will seem odd if the person rarely has token success.

⁶⁹ At least not for virtue ethicists. Consequentialists may have a different object of evaluation.

and avoiding false ones. In other words, epistemic exemplars are value pluralists. The remaining chapters give some sketch of what that type of inquirer looks like. In order to accomplish this, I need to make an important distinction between two kinds of epistemic virtue.

Two Kinds of Epistemic Virtues

Traditional virtue theories have identified an important type of virtue: those virtues that aim at the truth goal. Let us call these *veridical virtues*. Our ascriptions often identify virtues in agents because they aim at truth acquisition or error avoidance. For example, intellectual caution is recognised when agents attempt to avoid acquiring false beliefs. The agent who possesses this virtue knows when to be resistant to new beliefs or avoid poor inferences. Similarly, open-mindedness is a recognised virtue in which the agent seeks both to avoid error *and* acquire true beliefs at the same time. This virtue is practiced not only by having flexible beliefs, but also appreciating one's fallibility and limited cognitive power. Both caution and open-mindedness are veridical virtues; this kind of virtue is thought by many to be the only one that is relevant to epistemology proper.⁷⁰

I wish to distinguish veridical virtues from what I will call virtues of epistemic usefulness. The latter is rarely discussed or even acknowledged by

⁷⁰ See DePaul, Michael. 2001.

epistemologists. Virtues of epistemic usefulness are defined by their relation to non-veridical values. Employing this type of virtue may occasionally preclude the acquisition of truth; it may even require one to seek false beliefs. To be clear, virtues of epistemic usefulness do not *necessarily* preclude error avoidance and truth acquisition. Nevertheless, neither of those are primary aims.

What other value is there besides the veridical kind? My view is that there are many other kinds. While some of these values will emerge throughout the remaining chapters, I leave it open to as to how many there are. For now I classify them as epistemically useful to distinguish them from veridical virtues. I do not intend virtues of epistemic usefulness to be associated with practical value exclusively. These values are distinctly *epistemic*. I only use the term “useful” to denote the diversity of the values to which these virtues are aimed.⁷¹ That is not to say practical issues are irrelevant. It should be clear by the end that disposing with practical issues altogether is to abandon agent-based epistemology. This conclusion is irresistible once one recognises that agents are situated, embedded, embodied, and bounded. Nonetheless, the values at which these virtues are aimed are *epistemic*. Therefore, we can understand epistemic usefulness as a bundled term that picks out a cluster of values related by their *epistemic* nature. While the reader will get a clearer picture of these virtues and

⁷¹ I have resisted calling them “non-veridical virtues” because it gives the impression that virtues ought to be defined by their relation to truth.

values along the way, for now some understanding can be gained by contrasting them to the traditional veridical variety.

Here are a few candidates for these virtues: suitable self-deception, intentionally seeking false beliefs or avoiding true ones for the sake of understanding something or successfully navigating one's environment, knowing when to seek less than full knowledge, and knowing when to ignore low-level problems to address high-level ones even though solutions to the latter depend on those of the former. Interestingly, some virtues have a dual aspect; sometimes they are veridical, other times they are useful. They toggle back and forth depending on the way they are used by the agent.

For example, prioritising different levels of problems is aimed at the truth goal when employed to address macro problems such as scepticism; that is, when one ignores identifying the conditions by which one is able to discern that one is seeing a zebra.⁷² However, at other times the same virtue is aimed at values of epistemic usefulness. Other virtues have this dual aspect as well. While I will say more about these cases in the next chapter, I want to point out here that the existence of such virtues is evidence for pluralism. If one and the same virtue can be aimed at both veridical and non-veridical values, monism must be false. Moreover, there is no motivation to discount some virtue as essentially epistemic when the same intellectual skills are in one context aimed at veridical value while in another aimed at non-veridical value. They may have different

⁷² As opposed to another animal made to look like one.

token ends but they share the same *long-term* goal; and they employ the same skill. In other words, if virtue *x*—which is the employment of particular intellectual skills of judgement in environment *A*—is identified as an epistemic virtue with a particular token and type aim, then there is no reason to think *x* is any less epistemic if its type aim is the same while acknowledging its token aim is different. If the intellectual skills were epistemic in the former they must be in the latter.

My task is to give an account of these virtues such that they remain epistemic despite not being aimed at truth or error avoidance. This task is possible once we recognise that, like their moral counterparts, all intellectual virtues are aimed at two kinds of goals: immediate and long term—or token and type—goals. Epistemologists recognise the immediate truth goal; but they must also begin to acknowledge the long term goal of becoming specific types of inquirers. These long term goals are epistemic because they are centrally concerned with belief formation, belief maintenance, belief revision, belief transmission, epistemic autonomy and responsibility, as well as successfully navigating and tracking features of one's environment. Attempts to put a wedge between veridical virtues and virtues of epistemic usefulness will be unsuccessful. Efforts to classify the latter as non-epistemic create a cumbersome, unnecessarily fragmented, vacuous,⁷³ and ineffectual view of both the agent qua intellectual being and epistemology itself. More importantly, only by recognising

⁷³ Vacuous in the sense of purposeless.

the latter as properly epistemic is one able to account for a highly prized intellectual goal: wisdom.

The Telos of Virtuous Epistemic Agents

If the truth goal is the only epistemic good, we ought to think exemplary intellectual agents are ones with a high ratio of true beliefs. But this does not reflect our actual practices of epistemic evaluations, nor is monism reflected in real and widely recognised epistemic exemplars. It is telling that we identify intellectual giants apart from knowing their ratio of true beliefs. And we do not suddenly change our opinions about intellectual role models when we discover that many of their beliefs turned out to be false. Wayne Riggs notes that Aristotle has remained an exemplar despite the fact that a large portion of his beliefs were false.⁷⁴ We could compile a very long list of philosophers and scientists who were wrong about much of what they believed; and they would not be diminished in our minds for being members of that list. Thus, the epistemic ideal that is entailed by monism is incongruent with our actual practices. This disparity is due in part because philosophers have emphasised beliefs, as opposed to persons, as objects of evaluation. Once one makes *that* conceptual priority, the truth goal understandably takes the spotlight.

⁷⁴ Riggs, Wayne. 2003b.

There is value in evaluating beliefs. Sometimes we just want to know about the status of a belief, for example, whether it is true, false, misleading, or problematic. However, beliefs belong to people and are acquired by processes that often inform us about the beliefs themselves. This is an insight gained from reliabilism.

As has been noted by some philosophers, sympathy for reliabilism can lead one to have sympathy for virtue epistemology, where agents take conceptual priority as the objects of evaluation.⁷⁵ When agents move away from the periphery to the centre in theory-making, one ought to be reminded—even tethered to the view—that agents are anything but ideal. We do not have unlimited cognitive resources. We are not emotionless, objective, isolated human atoms bumping into other human individuals. Again, we are situated, embedded, embodied, and bounded.

My concern is to develop a view that is distinctly human-centred. Consequently, I conceptually prioritise *human* agents over other objects of evaluations, including beliefs and other species. This entails acknowledging the four aspects listed above. A virtue epistemology for angels, gods, eagles, and rodents will probably look substantially different. It is perhaps a limitation of the approach being put forth here that it will not be applicable to all species. I am not interested here in doing epistemology *in general*, for all beings. That is not to

⁷⁵ Many virtue epistemologists include reliability in their theories. Sosa, Zagzebski, Greco, Plantinga, and Goldman—among others—have shown varying degrees of sympathy for reliability and virtue epistemology. I should note that Plantinga has denied being a virtue epistemologist. However, his views appear to me quite friendly.

say that more general discussions of epistemology, for example analyses of knowledge and justification, are unhelpful and inappropriate. However, my aim is different. And I think there is substantial value in developing a distinctly human-focused epistemology.

If one were to give an account of divine virtue epistemology one would identify goals and practices that reflect supernatural cognitive capacities. It may be ideal, for example, to be infallibly logical, to avoid trivial false beliefs, and know even trivial truths.

Surely we must adjust our goals for those less fortunate. Virtues are, after all, not merely descriptions; they also have normative force. And normative utterances are meaningful only insofar as they are feasible. To illustrate, consider normative ethics. Some philosophers claim that "ought" implies "can". This is at least intuitively plausible; we think it is unacceptable to place a moral obligation on an agent who cannot possibly fulfill it. For example, we would reject a proposed moral principle that obligated human beings to save malfunctioning airplanes that are falling out of the sky and plummeting to the ground. The reason is quite obvious. It is not something that is within our grasp as human beings given our biological limitations.⁷⁶ Some argue that moral dilemmas also go beyond our obligations for the same reasons. I cannot, due to my physical limitations, both save the baby at one end of a burning building and

⁷⁶ Interestingly, we often do not recognise this of *all* beings. Many philosophers have pointed out that while humans could not possibly have that obligation, God does. This is what generates the problem of evil.

save the mother at the other end. So even though I may have an obligation to save each person had the other not also been there, I am not jointly obligated to save both. Regardless of what one makes of moral dilemmas, norms should not be in principle beyond the capacities of its intended agents.

Imagine that a group of surgeons go to a weekend workshop that promises higher surgical success rates if they follow new strategies. Losing fewer patients is something every doctor is interested in, so they are rather excited about their weekend. Now suppose the workshop strategy requires doctors to have a third hand and do computations in their head faster than any super computer can currently do. Surely our doctors are going to be less excited on the trip home from the workshop. The strategy may be perfectly successful if implemented; but it is clear this is not a strategy human surgeons can use. It goes beyond their physical capacities. Now, if the strategy consisted of using of new products and technology not requiring an extra limb and divine powers of computation, then not only would this be a good strategy, it might even be viewed as obligatory for *competent* surgeons. Suppose everyone agrees the latter strategy would not be as successful as the former one if they *could* both be implemented. The latter strategy would not be any less appropriate; nor would it make the former strategy any more appealing. The first strategy simply has no place; and we certainly cannot expect it to become commonplace or have some normative force in surgical theory and practice.

The goals, strategies, and norms of any human endeavor need to be appropriate. There is little point in coming up with a goal that is physically impossible or strategizing in a way that requires one to have the computational skills of an advanced computer. So when I claim I am interested in agent-based epistemology, it is tantamount to an interest in human-centred epistemology. It is not the epistemology of super humans, demi-gods, angels, or computers. If some theory in virtue epistemology fails to adequately acknowledge our finiteness and diversity, or assumes cognitive capacities that are out of our reach, then so much the worse for that theory. I recommend we take Timothy Williamson's advice and dismiss such accounts as simply irrelevant.⁷⁷ Why should we care about *that* kind of theory?

It is precisely because we are clever enough to recognise our own physical and cognitive limitations that we develop less than ideal strategies. That we use less ideal strategies is to our credit as intelligent beings. We may be stupid when compared to angels and gods, but what our epistemic heroes have done is exemplify how to be intelligent stupid creatures.⁷⁸ And I will argue that this is exactly what our ascriptions of epistemic virtues often pick out in agents we admire.

What then is an appropriate epistemic goal or telos for humans? I propose the telos is epistemic flourishing. That is not terribly informative, though

⁷⁷ Williamson rejects theories of knowledge that have become increasingly complex in order to avoid Gettier-type problems. As these theories get more complex, they become less likely to capture the kind of knowledge we have had in mind all these years. Williamson, Timothy. 2000.

⁷⁸ Adam Morton uses this expression.

we can note here that the expression is indicative of something long term. The telos is not fleeting or temporary, and so not reducible to token success—it is not simply a matter of acquiring a high ratio of true beliefs. Instead, it amounts to what Julia Annas calls the epistemically well-lived life.⁷⁹

Take for example two virtues commonly recognised in intelligent people. First, there is the virtue of knowing when to look for less than full knowledge. Second, there is the virtue of knowing how to prioritise domains of information.⁸⁰ Notice that neither virtue is part of traditional epistemology; the former virtue is often contrary! This should not be surprising to anyone. Much of our belief formation, belief maintenance, belief revision, inquiries, navigation through our environments, and problem solving are concerned with more than acquiring truths. And that is to our credit.

Imagine that a mathematician, John, is currently working on solving a problem that has plagued logicians for decades. John's best work is done while he walks around his neighbourhood. While John is crossing a street he starts to see the solution to his problem, causing him to temporarily stop walking. Unfortunately John is struck by a car and killed. John is clearly an intelligent person. However, epistemically speaking, this is not his finest moment, despite having come within seconds of solving a difficult math problem. While we can certainly acknowledge the high intelligence required to solve such a problem, we

⁷⁹ Annas, Julia. 2003.

⁸⁰ I will discuss these virtues in later chapters.

ought to withhold thinking John had epistemic success here. He exhibits quite clearly an epistemic failure. This failure consists of his inability to prioritise his cognitive resources. He should have known when to use his abstract inferential capacities and when to use his simple perception to acknowledge the danger of his present environment. His math skills should not have been employed to such a degree when crossing the street. He failed to prioritise correctly two domains, say, the logical and physical. Notice that John's case is about more than acquiring truth; it is about appreciating and prioritising the truths before him.

I suspect this is not a rare case. Agents often possess cognitive capacities but then lack the virtue of knowing when to employ them. Other times an agent's virtues may be thwarted by the presence of vices. Had John been more capable of multitasking he could have crossed the street safely while maintaining the computational power to solve the math problem. Given that multitasking diminishes the efficiency of the tasks being performed, he ought to have recognised his own limitations to do both tasks and compensated accordingly.⁸¹ This is generally true of daily life and problem solving. Failing to balance these intellectual tasks will not result in an epistemically well-lived life.

Take the following case that appears at first glance to elicit nothing more than moral success. Mother Teresa was known for her humanitarian efforts throughout her life and was a devoted Roman Catholic. Suppose, reasonably,

⁸¹ In chapter two I discuss the commonsense virtue of multitasking. I will argue that this virtue is not reducible to the truth goal.

that her humanitarian work was motivated by her belief in God. Suppose further that her belief was dogmatic because of her unwillingness to critically consider evidence against God's existence; she dismissed out of hand all potential defeaters. Despite her dogmatism, she was undoubtedly moral. Moreover, even if her moral character and behaviour were the result of a dogmatic and false belief in God, she would not be any less moral.

I think we can identify epistemic flourishing in the above case despite the overtly moral nature of her accomplishments. Suppose for a moment that there is some fact about the way the world is that entails objective moral value and a corresponding obligation to care for the poor. Mother Teresa's beliefs played a salient role in her successfully aiding hungry, poor, sick, and desperate people. Her beliefs (which may include a rich theology) informed and provided motivation and insight to sufficiently track her environment, pick out salient features, solve problems, and navigate them successfully. Surely we cannot disregard or diminish the role of underlying beliefs that directly motivate, inform, and guide an agent in successfully navigating her environment. Nor can we arbitrarily regard that success as non-epistemic. If belief x plays a salient role in problem solving or navigating agent A through some environment E , then some epistemic evaluation is appropriate.⁸² It just so happens this is a moral

⁸² A consequence of the view being presented here is that moral and epistemic successes are not mutually exclusive.

environment.⁸³ Interestingly, underlying beliefs do not need to be true in order for an agent to be successful.

Let us consider an explicitly epistemic case. Consider an agent in a maze who can reliably escape. Such a skill warrants making an epistemic evaluation. Suppose the agent in the maze adopts a strategy for escaping from the false assumption that she is in a computer simulation, similar to being a brain in a vat. She may infer from her background in computer science that computer programs will construct mazes using particular rules and predictably deceptive paths that lead to dead ends. Now suppose it turns out that she is not in a computer simulation, but nonetheless, in a maze designed by a computer. Thus, her general rules, even if subtle and not strictly enforced, reliably help her to navigate her way out. Our agent here is clearly exhibiting epistemic virtues and praiseworthy strategies. She is problem solving and using inferences in a virtuous way to achieve those ends. Now she might deserve higher epistemic credit if her strategy came from only true beliefs—because she fails to acquire token truths—but we would not withhold the epistemic aspect of our evaluation of her performance if we found out she received a monetary prize for escaping the maze or was motivated to escape in order to preserve her life. She can achieve both monetary and epistemic success.

⁸³ Remember, we have assumed that moral obligation is the result of some fact about the way the world is.

Similarly, a sea captain may find his position by interpreting data about the stars using a geocentric cosmology.⁸⁴ The method he uses is mistaken. But he succeeds nonetheless.

I see no relevant difference between Mother Teresa, the agent in the maze, and the sea captain. All of them succeed in navigating their respective environments by reflecting, inferring, and applying beliefs. Epistemic evaluations in all three cases are appropriate. And they all enjoy some level of success.

Perhaps one may try to drive a wedge between domains of success in the above cases. Such an effort however, will be awkward at best. We could not, for example, say that they used epistemic means to only non-epistemic ends. There are of course non-epistemic goals in the above cases. But these need not be mutually exclusive with epistemic ones. For example, it strikes me as bizarre that one must choose between moral and epistemic flourishing. If moral flourishing reflects some fact about the way the world is, then the two are not mutually exclusive. Epistemic goals are extensively and intimately intertwined in the variety of domains that describe the agent and her environment; one cannot simply isolate and pull at the epistemological thread without unraveling the agent's proper functioning in the world. On the contrary, properly functioning agents *intentionally* use their epistemic capacities in such ways that intersect many domains. And when they are appropriately used, they are successful, not only morally or practically, but often also intellectually. They deserve *intellectual*

⁸⁴ Adam Morton has used this example.

praise for integrating their beliefs in the variety of environments in which they find themselves.

Since my aim is to reconsider the parochial nature of epistemology, it amounts to questioning begging to disqualify out of hand any success that is not reducible to truth acquisition and error avoidance.

One reason I suspect we do not typically think of cases such as Mother Teresa's as epistemic flourishing is that we tend to acknowledge the independence or primacy of the moral domain. Her moral merits overshadow her epistemic ones. We should not think, however, that her epistemic accomplishments are non-existent. They may live in the shadow of her moral virtues, but if helping the poor is an obligation that follows from the way the world is, then her recognition and appreciation of this fact reflect an understanding of the world.

The conclusion to draw from the above is that the telos of virtuous epistemic agents goes beyond the veridical one. I have done little more than cash out epistemic flourishing into what it is not. Here I only offer a brief sketch; a complete account of the epistemically well-lived life depends on more underlying metaphysics than I wish to assume. For example, it will depend on whether or not God exists. If He does, then it will be informed by a divine will. The status of God's existence will inform us as to what kinds of character one ought to be aiming. God may require that people obtain a sort of humility not otherwise associated with epistemic exemplars. Or He may require one to adopt

some kind of epistemic resignation with regard to curiosity and inquiry. Much of that will be entailed by the character of God. On the other hand, if God does not exist, that will also inform us as to what constitutes the epistemically well-lived life.

Without taking too much for granted about the way the world is, we can identify a few features of the epistemically well-lived life. Consider John, the mathematician above. Suppose instead of being hit by the car John suddenly broke his concentration and focused on the oncoming traffic; he then jumped out of the way to save his life while losing the answer to the math problem. We should think he acted, to his credit, with commonsense. Moreover, he avoided a couple of failures, losing his life being the biggest one. That he failed to arrive at an answer to the math problem is trivial by comparison. In fact, I propose that his success is at least partly intellectual since he knew how to navigate his environment. He tracked features of his environment, picked out salient information, and prioritised them correctly. One may point out that John sacrificed one token (math problem) for the sake of another (the belief that oncoming traffic is life-threatening). However, he also contributed to a longer term goal of becoming and maintaining a *type* of epistemic character. This character includes acquiring traits that are conducive to being an efficient and reliable inquirer and problem solver, where problem solving is widely employed. Cultivating particular intellectual characters is the one feature of the telos.

Our goals as epistemic agents ought to go beyond individual desires. Our personal interests surely play an important role, but once one acknowledges the need to index virtues to the agent's environment, one can see immediately that the epistemically well-lived life is informed by some corporate or communal aspect. Our world includes among other things, communities. So a full account of the telos will need to have a central place for groups and communities.

Many details must be left until one has a developed metaphysics. For that reason, it should be clearer now why I find Aristotle's "eudaimonia" helpful. Eudaimonia is not domain specific. The diversity, relevance, and longevity of epistemic virtues require a telos that is multilayered and reflects the agent and the world in which she finds herself. It cannot be reduced to token moments of acquiring true beliefs and avoiding false ones since it includes the development of particular types of intellectual skills that make up the agent's life and character. Consequently, not all agents should be aiming for the same epistemic character. In addition, diverse intellectual skills are often more advantageous for communities to navigate their complex environments, which include, not only natural but, socially constructed ones as well. Therefore, one should be suspicious of the view that the *epistemically ideal* looks the same from agent to agent.

It should be clear that this departs quite sharply from the traditional truth goal. The epistemic aim is widened to include other values. Recall Chisholm's early view:

We may assume that every person is subject to a purely intellectual requirement—that of trying his best to bring it about that, for every proposition h that he considers, he accepts h if and only if h is true. One might say that this is the person's responsibility or duty qua intellectual being.⁸⁵

Interestingly, Chisholm abandoned this view and later wrote:

I have previously written, incautiously, that one's primary intellectual duties are to acquire truth and to avoid error. What I should have said is that one's primary intellectual duties are to believe reasonably and to avoid believing unreasonably.⁸⁶

Chisholm appeared to think rational belief formation was valuable in addition to the truth goal. Insofar as the former is not reducible to the latter, it is a move away from monism.⁸⁷ While, I am making no claims here that Chisholm intended to endorse full-scale pluralism, this is certainly a departure from his earlier parochial description of epistemic exemplars. Likewise, the remaining chapters attempt to reframe epistemology in more than veridical terms and value.

⁸⁵ Chisholm, Roderick. 1986. 47-54.

⁸⁶ Chisholm, Roderick. 1989. 1.

⁸⁷ Kvanvig takes Chisholm here to be abandoning monism. Kvanvig, Jonathan. 2003.53.

Chapter 2

Commonsense

Commonsense has a long history in philosophy. It has enjoyed both its advocates⁸⁸ and critics.⁸⁹ It has been employed in various ways in the literature. It sometimes denotes a skill held by people. Other times it is treated as a property of beliefs. Some describe it as a disposition or stance one takes toward some topic or life in general. On the one hand, it has been a term of praise. On the other hand, it can be a criticism when it implies naivety in a particularly unique context. Because commonsense has a variety of meanings and uses, I need to be clear what sense of the term I intend here. My discussion of commonsense will be couched in terms of epistemic virtues. This is not to say it cannot be examined in other terms, but ascriptions of commonsense often do denote virtues possessed and employed by the agent. Some of these virtues are particularly relevant since they imply pluralism.

As I use the term, commonsense is not any one capacity, skill, or virtue; there is no virtue of commonsense. Instead, I take it to be to be a collection or set of related virtues. I do not have much interest in describing necessary and sufficient conditions for what makes a virtue of the commonsense variety, if there are any such conditions. These virtues may share nothing more in common

⁸⁸ For example, G.E. Moore and Thomas Reid.

⁸⁹ For example, David Hume and sceptics in general.

than that we tend to think a sensible, down to earth, or minimally competent person ought to possess such intellectual powers. We probably do not have names for many of the commonsense virtues. Nonetheless, one can identify them with commonsense once one appreciates the intellectual skills behind some performances.

Commonsense ought to be distinguished from good sense. Commonsense is not always a term of praise.⁹⁰ There are environments in which commonsense qualities and procedures are inappropriate. We may, for example, criticise someone for using commonsense interpretations of particular genres of literature, film, and art such as mythology, fantasy, science fiction, and apocalyptic. Or we may wish to entertain a thought experiment that requires one to leave commonsense aside. Thus, virtues of commonsense—like other assets and skills—are context-sensitive. Treating commonsense as a set of virtues captures this feature since when either one is used in the wrong environment, we think the person is using a skill inappropriately.⁹¹

Identifying commonsense with a set of virtues also has the advantage of explaining why ascriptions of commonsense vary a great deal. For example, it is commonsense to shop around and compare bank quotes when securing a mortgage; it would show a lack of commonsense to go with the first offer. Likewise, a person is practicing commonsense if she brings an umbrella with her

⁹⁰ Unlike good sense.

⁹¹ On the wrong occasion.

on a walk when it is raining. Both people display commonsense by recognising particular relations. The hopeful home owner is counting on competition between banks to generate a better interest rate. The agent going for the walk recognises a logical relation between the following:

- A. It is raining outside.
- B. If it is raining outside I will get wet when I go outside.
- C. But, I do not want to get wet when I go outside.
- D. If I bring an umbrella, I will not get wet.
- E. Therefore, I should bring an umbrella.

A third agent may practice commonsense by taking the reports in her newspaper at face value, as opposed to employing some aesthetic hermeneutics that is more appropriate for aesthetic literature. This third agent appreciates the genre of her literature and is taking up a commonsense stance to gather information. Had she picked up a piece of fiction that was written in the style of a news report, she would need to employ a different hermeneutical approach.⁹²

Sometimes commonsense consists of recognising a particular relation or pattern between beliefs, however, not necessarily the same relation or pattern every time; and sometimes it purposely ignores relations and patterns.

⁹² This hermeneutical error occurred in 1938 when a radio show "The War of the Worlds" was given in the style of a news bulletin. Some listeners panicked, thinking an alien invasion was in progress since they did not realise it was a radio show, not a news report.

The following virtues can be identified with the commonsense variety. I will narrow my attention to virtues that do not aim at the truth goal. Nonetheless, they are practiced by virtuous agents, contribute to the epistemically well-lived life, and are quite common.⁹³

Knowing when one has Enough Evidence⁹⁴

Knowing when one has enough evidence, knowing that more evidence is superfluous, is a commonsense virtue. One might immediately recognise this virtue in researchers and scholars. However, this is commonly used even among moderately intelligent people. In fact, it is striking when we notice someone lacking this ability; and it might cause us to shake our heads and wonder what that person was thinking, or *if* that person was thinking.

Take the following case. Suppose Jones and his wife return from a three day vacation. Jones enters their apartment and smells what he takes to be natural gas. He turns to his wife—who always seems to have a better sense of smell than he does—and asks if she smells anything. She says she smells nothing but stale air. So Jones thinks it must be nothing. He and his wife go out for dinner that night and return two hours later. Upon re-entering his apartment, once again Jones smells what he takes to be some sort of gas. He turns to his wife and

⁹³ Common, but not always present. As mentioned above, there is a virtue of knowing when commonsense is not good sense—for example, in a house of illusions or a substantially different cultural setting.

⁹⁴ This may be a collection of virtues such as: knowing when to stop looking for evidence, knowing when to start looking for evidence, and knowing when not to look for *any* evidence.

asks her if she can smell any gas. Once more she says that she smells nothing. Jones passes it off as nothing. For the next several days Jones encounters the same gas-like smell whenever he comes home but does nothing more than pass it off as some sort of sensory error; he does not even bother asking his wife anymore. However, the next day Jones hears a knock at the door. His neighbour from the apartment above him stands at the door asking him if he smells any sort of gas. Shocked, Jones says he has smelt it for days. His neighbour tells him that another tenant has also been smelling gas. His neighbour decides to call the gas company to come see if there are any gas leaks. It turns out there is a gas leak; and the leak is coming from Jones' apartment. The leak has spread so much as to be detectable from tenants above him.

Jones and his wife lack commonsense. They have evidence for thinking there is a gas leak, namely, Jones perceiving gas. Of course, they both have some evidence for thinking there is no gas leak, namely, Jones's wife smells nothing. Historically she has a good sense of smell. However, both of them should have sought out more evidence. They could have asked a few neighbours if they smelled gas. Or more cautiously, Jones and his wife could have skipped the testimony of neighbours altogether and called the gas company to test for a gas leak.

Their lack of commonsense might anger his neighbours, "You smelled gas and did not think it was sensible to investigate the potential danger to you, your wife, and the entire apartment building?" They may verbally reprimand Jones

and tell him that it is just plain commonsense to take every precaution when dealing with gas.

Jones did not know when he had enough evidence. He should have asked several people or called the gas company immediately. Jones could have lacked this virtue in another way. Suppose after his wife told him the first time that she did not smell gas he decided to call the gas company anyway. Now suppose the gas company arrived, performed tests, and concluded what Jones suspected; there is a gas leak. He then asks his neighbours if they smell gas. They tell him they do but are aware that the gas company has sent someone out to fix the problem. After talking with his neighbours, Jones goes to the library to take out several books on gas leaks and how to detect them in homes using everyday household items. He then researches gas leaks on a few websites for several hours in an attempt to determine the likelihood that there really is a gas leak in his apartment. He concludes there is a gas leak.

In this alternative case, Jones lacks the virtue of knowing when he has enough evidence. And again, his neighbours are displaying it. He lacks the virtue because he looks for much more evidence than he needs. In fact, he probably already has in his possession the strongest evidence he is going to acquire; the gas company's instruments have detected a gas leak. Continuing to look for evidence, which probably will not conflict with or be as reliable as the evidence he already possesses, is not sensible. His neighbours' testimonies are not as reliable or as strong of evidence as the gas company's instruments. Likewise, the

evidence gathered from books and websites on generic gas leaks are too weak compared to the evidence he already has. So they will do little to counter or reinforce the evidence provided by the gas company. If so, then there is little point in looking for more evidence. In this case, Jones's lack of commonsense is so grossly evident that we make think there is something particularly wrong with him. We may opt for a stronger ascription in other cases when people look for far more evidence than one ought to: paranoia or even obsessive compulsion.

In the first case Jones stops looking for evidence too early. In the second one he goes too far. In both scenarios his neighbours seem to know when they have enough evidence. This virtue, like others, is context-sensitive. While there is little hope of giving an exhaustive list, we can make some general observations about some salient conditions that people need to consider when employing this virtue. These conditions are informed by the situatedness, embeddedness, embodiment, and boundedness of the agent.

First, knowing when one has enough evidence depends on what is at stake. Generally, this means the ratio of resources spent to the potential payoff and cost. Let us first consider resources spent.

Looking for evidence takes cognitive resources, physical energy, time spent, sometimes money, favours, and more. Each resource used in securing

evidence possesses a particular value;⁹⁵ that value is mostly dependant on the agent but also on other environmental factors.

For example, whether it is worth it for an agent to spend one hour of time gathering evidence or problem solving will depend on whether one's time could be better spent elsewhere; economists call this "opportunity cost". The time it takes to gather evidence will vary due to cognitive or physical constraints of getting that evidence;⁹⁶ and people need to judge if their time could have been spent, for example, accomplishing two or more other tasks. Maybe the evidence gathering is taking up too much of the agent's attention so that she is too fatigued to make impending inquires that are more important. Of course, what is mentally fatiguing to one agent might not be for another. So this is a variable to be determined on a case to case basis.

The value of a resource will depend on how much of it the agent has, for example, how much time one has during which a particular task can be done or indeed one's life expectancy. So not only will the physical and cognitive abilities of the agent⁹⁷ be relevant, but so too will the physical conditions of the environment.

Evidence gathering can come cheaply or it can be expensive. That often depends on who is doing the inquiring. However, the potential payoff and

⁹⁵ This value is instrumental and not exclusively epistemic.

⁹⁶ Perhaps the evidence is in an inconvenient location or at a great distance.

⁹⁷ Of course, individuals are not the only ones looking for evidence. Communities, countries, and corporations all look for evidence at the cost of resources. So these considerations also apply to groups.

likelihood of that payoff also figures into whether those resources are worth spending. Searching for more evidence may take enormous resources. Yet it may be worth it given the potential rewards. Let us consider monetary payoffs. Corporations often invest millions of dollars in research and development hoping to yield more in return. Thus the motto, "You have to spend money to make money". Somewhat similarly, corporations often invest because they "cannot afford not to". This expresses the negative consequences of not investing or searching for evidence of potential markets. Perhaps other corporations are growing in size and threaten the stability of current profits. Thus, the ratio being considered is resources spent to both net gains and losses; and this analogy carries into the epistemic domain. Sometimes the agent cannot afford the potential loss of not looking for evidence, despite the expensive resources required to secure that evidence.

In addition, the amount of resources the virtuous agent uses depends on the *likelihood* of payoffs. Potential payoffs are just that, potential; they are not guaranteed. Their probability may be high, low, or completely unknown.⁹⁸ Thus, the virtuous agent may decide it is not worth spending resources on a potentially large payoff if she recognises the likelihood of the payoff is quite low. Take, for example, lotteries. One can buy lottery tickets with little resources—a few dollars—for the potential payoff of millions of dollars. Yet, many people do not

⁹⁸ We may start research or evidence gathering with little to no idea of potential payoffs; and it may be virtuous to do so.

think it is worth buying lottery tickets as an investment. Perhaps lotteries are worth it for the fun of “playing the lottery”. However, they are not a sensible way to invest resources in order to make a profit. In fact, we think it is rather foolish of someone who is convinced they are good investments. There are epistemic equivalents to the lottery case.

One candidate might be something like Pascal’s Wager: If you believe in a particular God and he really exists, when you die you will receive eternal life and happiness. If you do not believe he exists and he really does, you will be severely punished. You have these two options presented to you. However, the person presenting the choice points out one more piece of information. If you believe and it turns out he does not exist, you have lost very little—relatively speaking. So, the investment to believe or worship such a being is rather little considering the potential reward and safeguard against great cost.

Still, many find it not worth the investment, at least not based on those particular grounds. One reason is that many think that the probability of the payoff is quite low. Those who think the investment is worth it, usually have other reasons for believing God exists.⁹⁹ Thus, we have a candidate for sensibly denying even a small investment of cognitive resources for a potentially great payoff and avoidance of negative consequences.

⁹⁹ I think the wager presented here is probably not what Pascal intended. Since Pascal gives other reasons—for example, the explanatory power of Christianity—he did not intend his wager to stand alone as a simple choice between a potential eternal reward and punishment.

The decision not to inquire is much more frequent than one might expect. Just as our decision maker above may decide the investigation into God is not worth her resources, people without an aptitude for particular areas of science, math, language studies, and so on, may decide some of those areas are not worth the resources required to learn them. For example, if Mary knows there is very little chance she will understand high level calculus or quantum field theory, then she is likely to decline any opportunity to inquire into these fields of study.

We can now identify one reason Jones lacks the virtue of knowing when he has enough evidence. Given what was at stake, he should have continued to look for evidence beyond his own sensory perception and his wife's conflicting testimony.¹⁰⁰ When we consider the potential cost was his life, his wife's life, and the safety of the entire apartment building, as well as the little cost of resources involved, he should have furthered his inquiry by asking his neighbours about the odour and called the gas company.

Many of these considerations are irrelevant when epistemology is done in a vacuum. They are not even recognised as relevant *epistemic* factors by many philosophers. But why not? They figure into *all* of our epistemic practices, including our evaluations. We know we do not live in vacuums. We have to make our way through life, solve not just one isolated problem, accomplish one atomized inquiry, and acquire more evidence for discrete beliefs, but rather do

¹⁰⁰ This involves the virtue of knowing what kinds of investigation to carry out.

so against a background of dynamic environments. One mistake epistemologists have made is assuming that because something has practical, moral, or physical value, that it is a practical, moral, or physical resource, it cannot also be an epistemic one. The epistemic domain is not mutually exclusive with other domains of life.¹⁰¹

A second factor associated with this virtue is the quality of evidence one already possesses. Our beliefs often come with some degree of evidence. It may be perception, data from scientific experiments, testimony from any number of sources—media, friends, strangers, computers, and instruments—or it may even be intuition. A person may have no evidence other than that she already has a confident belief though forgot how she came to acquire it. These sources of evidence vary in degree of strength and reliability. We might, for example, regard the testimony of a friend or expert as more valuable than our own. Couples exemplify this when a husband asks his partner what to order at a restaurant, whether he has seen a particular movie, or when he seeks guidance concerning his life. We often recognise that other people have superior perceptual skills, a vantage point, or authority, even about *our* lives.¹⁰² That we

¹⁰¹ Adam Morton pointed out to me that very little attention in epistemology is spent on the design of experiments and guiding what one should *do* to get different evidence. Instead, the attention is on what one should think given the evidence one already has. Notice that once one attempts to give guidance on *doing*, one must take into consideration the variety of limitations that people have. A good guide will not advise someone to do something they cannot do.

¹⁰² For example, counselors, professionals, and those better educated about the object of the inquiry.

take some evidence to be stronger or more valuable than others is obvious.¹⁰³ Since inquiries usually begin with some evidence already in hand, one needs to determine the value of that evidence. This factor should strike the reader as being properly epistemic.

I want to mention here a rather interesting consequence of knowing when to stop looking for evidence when we consider very high stake cases. The virtuous person will realise that on the one hand sometimes even redundant evidence is “worth it”. On the other hand sometimes inadequate evidence suffices. Moreover, both can be true at the same time.

Redundant evidence is sometimes given value by the agent, even while acknowledging its redundancy. For example, one may check and recheck whether a dangerous appliance was left on before leaving the house. People often double check before crossing the street “just to be sure” even though they were sure before rechecking. One may double check the instructions on a prescription each time if the drug being taken is a particularly potent one. Or perhaps a young woman is driving in an infamously dangerous section of the city. She may double check to see if her doors are locked, even though she knows they automatically lock when the ignition is started. She may have even bought the car for that safety feature. However, given that the neighbourhood she is driving through is notoriously dangerous for women, she feels the need to check if the doors are indeed locked. None of the above cases of redundant evidence

¹⁰³ For a variety of reasons, which we can set aside here.

appears to be a waste of resources. In each case it seems virtuous to look for redundant evidence.

Given high enough stakes, we sensibly take precautions. Often this takes the form of seeking out redundant evidence. This kind of precautionary evidence is akin to two faithful lovers who have been tested for and cleared of STIs, and yet still wear a condom during sexual intercourse. The two are well aware of the test results, yet they wear protection because the stakes are high. While the use of a condom is redundant—since they have been tested—we would hardly call the practice useless or foolish on their part. The virtuous nature of securing redundant evidence reflects the agent's appreciation of human fallibility. In fact, we may criticise someone for failing to seek redundant information just as we might if the two lovers do not use protection.

Alternatively, sometimes it is virtuous to act on inadequate, poor, or unreliable evidence. Take the case of Mary and her brother Larry. Larry has a history of making false claims, often for the purpose of playing tricks on Mary. He has been doing this since they were children and has continued the practice into adulthood. He often makes outrageous claims. For example, when he was a teenager, he convinced Mary that their parents died in a car accident. This sort of lie is typical of Larry; he is the family joker, equally unappreciated and unethical. Suppose while at work, Mary gets a phone call from Larry. Larry claims he has just been over to Mary's house to borrow her vacuum and while there discovered her husband had apparently slipped off a ladder and is bleeding quite

a lot from his head. Without hesitation, Mary hangs up the phone, dials 911 and jumps in her car. She drives home to discover there has been no accident and her brother Larry has pulled yet another prank on her. Impressed with himself, Larry asks, "Why would you believe that? You don't even own a ladder!"

Given her brother's history of these types of unethical lies, and that she does not own a ladder, Mary had at best inadequate evidence to act the way she did. However, her actions seem perfectly sensible. When the stakes are *that* high, inadequate evidence—including inadequate reasons—can be sufficient. This seems especially true given limited resources. In this case, time was a limited resource.

Take the following case that is quite common. Betty is walking home from work at night. She has several routes she may take to get home. She usually takes the same route home as it is the quickest. One night however, while walking her usual route, she notices a man walking down the same street. The man *looks* rather respectable; he is wearing a nice suit and seems to be minding his own business. But she has a bad intuition about him. She does not feel safe being on the same street as him at night. If asked, she could not pick out what makes her feel unsafe; but she cannot resist the feeling that the longer she stays on the street with him, the more she is in danger. Betty can turn down any number of streets to take a different route. So she takes an alternate route leaving the man to continue walking on his own.

We often have bad intuitions about people and situations. Like Betty one can have an intuition that a person is untrustworthy or suspicious; one can have the intuition that one's current environment is not safe, beneficial, or ideal. Betty does not appear to have any good reason or evidence for taking a different route. All she has is an intuition about the well-dressed stranger. Nevertheless, one can reasonably sympathise with her actions, even praise her for being a sensible woman. We would hardly say she has adequate reasons or evidence to take an alternate route. However when we consider the stakes, we can see that adequate reasons and evidence are not required. In fact, they are not desired. It would be risky to seek adequate evidence before she took another route. Waiting around to see if the man displays an adequate sign of threat is hardly the virtuous thing to do; and I suspect no sensible person would endorse such a practice. The sensible thing to do is to take another route *before* the threat manifests itself. Betty knows when to stop looking for evidence, which entails *not* knowing whether the man presents a real threat to her.

Now, one might object by claiming that it makes *practical sense* to act before she has evidence, but not intellectual sense. This is the sort of fragmenting I think is mistaken. There is no meaningful sense in which one can carve off Betty's intellectual performance. Her intellectual performance is constituted by the inference and consequent decision to change routes based on her environment. Clearly this is a practical inference; the payoff is a matter of

personal health and safety. But that does not mean it is any less intellectually praiseworthy. We *would* praise her intellectual character and skills.

A different objection grants that Betty was reasonable to take another route and reasonable to conclude that the man was likely enough to be a threat to warrant the action, but denies that she is warranted in *believing* he was dangerous. Nevertheless, her inference to the possibility of danger remains epistemic.

These types of considerations are necessary because we need to get things done in the world. Often we must form a belief or make a decision based on inadequate evidence. If I am approaching a crossroad and have two available routes to work that are equal distances, I need to decide which route is going to be quicker based on an observation about the density of traffic in each of their respective directions. It makes little sense to wait for adequate evidence. I need to form some belief and decide.¹⁰⁴

It should be clear now that knowing when one has enough evidence is a commonsense virtue that is highly contextual and often aimed at the truth goal. However, when we consider the variety of contexts and uses of this virtue, it becomes more difficult to ignore the plurality of value to which this virtue is aimed.

¹⁰⁴ Or perhaps I need to form a view about potential employees based on interviews that provide me with inadequate evidence. If I need to hire one of the candidates, I need to make a decision one way or another.

Consider again the above cases. Lacking this virtue hindered Jones in attaining the true belief that there was a gas leak. Instead he acquired a false belief. Since these are traditional epistemic values that are being aimed at, knowing when one has enough evidence is indisputably an *epistemic* virtue.

However, the aim of this virtue is not always so traditional. What was the cost and benefit of having the virtue in the second case where Jones continued to look for evidence when he should not have? It was not that he might have obtained a false belief or failed to gain a true one.¹⁰⁵ Instead, lacking this virtue caused him to waste intellectual resources.¹⁰⁶ Our observation that he lacked commonsense was not due to some failure to achieve the truth goal. It is not simply that we think he should have been using his resources to find other truths; he would have a higher epistemic standing—he would have the virtue—just by doing nothing. So this virtue is not always aimed at the truth goal.

Consider the cases of Mary's lying brother, Betty's intuition, and driving in dangerous neighbourhoods. All of those agents knew when they had enough evidence, but the virtue in these cases is not aimed at true and false beliefs. Nonetheless, they all employ the same virtue that Jones lacks in the first case, which has been identified as an epistemic virtue. Therefore we have a commonsense epistemic virtue that acts in some contexts as a veridical virtue and in others like a virtue of epistemic usefulness.

¹⁰⁵ That may be a possibility, but it is not the reason for thinking he lacks this virtue.

¹⁰⁶ Including other resources such as physical energy, time, and perhaps money.

Those wishing to resist such a conclusion may be suspicious of a bait and switch, that while Jones exhibits this epistemic virtue, the other agents are really displaying a different virtue, only with the same name. This, however, is unlikely.

All of the people above are displaying skilled judgements about when they have enough reasons or evidence for having a particular belief or taking an appropriate action. This consists of recognising the quality of evidence one has and being able to evaluate it in light of potential payoffs and costs. In some cases, virtuous agents know when to look for more evidence, when to look for more reliable evidence, when to look for redundant evidence, when to stop looking for evidence, and when to look for no evidence other than the minimal, weak, and unreliable evidence they have already obtained. The stakes can change the aim without changing the skills being employed.

Thus, there appears to be no relevant epistemic difference between these agents. They appear to be using roughly the same intellectual capacities skillfully. Now, if the critic points to the different aims of these virtues as being the epistemically relevant reason for thinking they are not the same virtue, she is question begging. That is the very thing in dispute. It would be like saying that eyesight is a different capacity when it is used in bird identification than when it is used in threading a needle. Thus the critic cannot dismiss these cases as different virtues just by the difference in their aims. The critic will need to pick out a difference between these cases that is both *epistemically* relevant and non-

question begging. Once one recognises the cases above employ the same intellectual skills, pluralism is implied.

Consider the alternative, that different aims entail different virtues. It is more likely, I think, that virtues can vary in token aims depending on context, even if virtues of the same *kind* have aims of the same *kind*.¹⁰⁷ For example, a moral virtue often has a diversity of token aims in different environments even though all the aims are of the same kind: moral. Likewise, epistemic virtues have a variety of token aims in different contexts, but those aims are still epistemic in nature.¹⁰⁸ So while they may not aim at token truths, they still contribute to the epistemic character that the virtuous agent develops to successfully navigating her environment. They figure into the epistemically well-lived life. If knowing when one has enough evidence is identified as an epistemic virtue in one case, then surely when those same skills are employed in another context they are also epistemic despite having different token aims. If so, our epistemic evaluations of praise include a commonsense virtue that cannot be accounted for by value monists.

¹⁰⁷ Moral courage, for example, may have more than one aim. It is true that all those aims will be moral in nature, but there is no reason to think each virtue has one token aim. If so, the list of virtues would be so numerous as to be useless.

¹⁰⁸ I do not mean to suggest here that these aims are exclusive in nature. I suspect they are often not so easily distinguished. Some virtues or set of virtues may work together to aim at both moral and epistemic goals.

Knowing when one is in over one's head

The second commonsense virtue I want to identify is knowing when one is in over one's head. Alternatively, one can describe this virtue as knowing one's limit of competence. This is a necessary virtue in order to employ a further one: knowing *what to do* when one is in over one's head.

This virtue is manifested in many common situations. For example, students often exhibit this virtue when they seek alternative university courses from the ones in which they are currently enrolled. Many students find, for example, math courses intellectually overwhelming. However, once an agent is aware of being in over her head she has an advantage over those who are unaware. A student not having this virtue may receive a poor grade in the course, experience increased stress and frustration, find that the course consumes more time than she has available, or even fail the course.

An agent who has this virtue is more capable of practicing other epistemic virtues such as caution, managing intellectual resources, and knowing what to do when one is in over one's head. Caution might include rejecting or filtering out some of the course information that is likely to mislead or confuse the student. Managing one's resources in such situations might include securing additional resources, for example, a tutor. Or it may consist of devoting a great deal more intellectual powers to the subject.

Suppose Sally has saved a math course for a semester in which she has few demands elsewhere in her life since she knows it will be intellectually demanding. Knowing how difficult math is for her, she may even decide to seek the help of a tutor. It is to her credit that she recognises her limitations and compensates by seeking help from another agent, presumably one who not only has an aptitude for math, but also the skill to teach it. Interestingly, Sally's goal in seeking help may not be to obtain a high ratio of true beliefs. Her goal will depend on several factors. Perhaps she may not need an A in the class. Or maybe she is hopelessly terrible at math and can only reasonably be tutored enough to get a passing grade. If so, putting in the effort to achieve such a high degree of true beliefs is futile. There is nothing extraordinary or intellectually blameworthy about this strategy. On the contrary, she deserves credit for managing her intellectual limitations. In fact, we justifiably criticise people for reaching well beyond both their physical and intellectual limits. Of course, if she is well in over her head and knows it, Sally may drop the class entirely and avoid acquiring any true beliefs whatsoever.

None of those choices is epistemically perfect of course. However, Sally is not an ideal agent. She is terrible at math. What good then would it be to evaluate her only by an idealised epistemology with one goal? One would end up giving terrible advice to Sally under such a view. Sally has no hope of achieving ideal ends. Better that she acknowledges her intellectual limitations and adopt a strategy that reflects her *real* capacities and situation. Moreover, pluralists can

agree that acquiring more true beliefs¹⁰⁹ is epistemically better. What they deny is that this is the *only* way in which she can achieve an epistemically praiseworthy status. If she drops the course because math is hopelessly beyond her cognitive capacity, she deserves praise for adopting such a strategy despite not acquiring more true beliefs.

The monist might be tempted to say that dropping the course is nothing more than a practical or academic strategy.¹¹⁰ Indeed it may be both. But that does not mean it is not also the intellectually virtuous one. She may fail to acquire some token true beliefs in math, but dropping the course might contribute to or reflect an intellectual proficiency that is particular or suited to her mental powers. Some intellectual personalities tend to exclude others. One reason is limited resources. Specialising in or developing an area means less time for other areas. Another reason is due to facts about human psychology. Creative personalities might be less detail oriented. Disciplined people might find it more difficult to be sympathetic. Once virtues are understood to contribute to the agent's character and intellectual *life*, counting token true and false beliefs ceases to be an exhaustive analysis.¹¹¹

One further point is worth mentioning. Being in over one's head is not

¹⁰⁹ And avoiding false ones.

¹¹⁰ She may also cheat if she realises she is in over her head. My personal view is that this is almost always a case of failure, though I would not dismiss out of hand the possibility that some environments warrant such behaviour.

¹¹¹ She may even receive an A in the course and still obtain an unfavourable epistemic evaluation. Suppose she received an A in the course only by compromising other obligations, including intellectual ones. In this case, even though she had academic success, she failed to know what to do in such a situation—she spent too many resources on the course such that other areas suffered.

always something virtuous agents avoid.¹¹² Sometimes one must go beyond one's competency limit to gain some information that could not otherwise be obtained. This might seem friendly to monism. However, on closer inspection it amounts to a counter-example. When virtuous agents put themselves in situations beyond their limit of competence to acquire information, it is often because they are seeking something other than true beliefs; they are looking for beliefs sufficient enough to act on, solve problems in a related area, get a sense of the field or lie of the land.

Knowing when to look for Empirically Adequate Beliefs

Sometimes we purposely put ourselves in situations we know will end in incomplete knowledge. We do this because we recognise that sometimes true beliefs and full knowledge are either not necessary or not desirable. This is why Sally may only seek a passing grade in math. Sometimes knowledge is superfluous knowledge. Other times the virtuous agent recognises the difficulty

¹¹² A moral example might be the following. While waiting for the fire department, Odie sees a child in a window of a burning building. Odie is well aware of just how much the fire has spread. And he is aware that it is highly unlikely that he will reach the child in time. Yet Odie enters the building anyway knowing full well that he is in over his head. The outcome to this seems to me irrelevant. It does not matter whether Odie succeeds or fails to rescue the child. He might be courageous even if the child is not saved. The effort itself is praiseworthy and contributes to his moral character.

involved in obtaining complete knowledge or even a high ratio of true beliefs and correctly decides against such a goal.

This is commonly used when multitasking, a distinctly epistemic capacity. Virtuous agents know how and when to multitask. But multitasking involves splitting one's attention, often sacrificing full knowledge, understanding, and true beliefs of one topic for the sake of incomplete, fragmented, or empirically adequate beliefs of three or four topics. In fact, it would be epistemically defective of an agent if she did not multitask from time to time.

An agent may seek only empirically adequate beliefs for reasons unrelated to multitasking. It may be that a high degree of knowledge, understanding, and true beliefs simply provide no advantage for the agent in her current situation. For example, perhaps Jones only wants an approximation of the temperature of the city to which he is travelling; he wants to dress for the weather. But he does not need to know or understand the forecast to accomplish that. The information he gets does not even need to be true. When he asks his assistant about the weather, she may simply guess that it is twelve degrees. Or she may use a more ambiguous description, perhaps telling him that it is warm. If he wanted he could seek more precise and reliable true beliefs. But it is not clear it is always worth his effort.

Moreover, the *kind* of information may be such that seeking more of it is inappropriate. Perhaps Jones has stumbled upon some office rumours that co-workers have spread as gossip. Suppose the content of the gossip is sensitive or

private. It would not be epistemically virtuous of Jones to seek clarification or full knowledge of these rumours. On the contrary, it would be a misuse of his epistemic resources. Part of the reason for this is that he would be failing to identify and appreciate the kind of information he has received. We identify this as the vice “nosiness”.¹¹³

Some information is valuable but private. Solving a problem may require knowing private or intimate information but not details. For example, a physician diagnosing an STI may require a history that includes sensitive information, but not details. It *may* require details; but a good physician will discern when to look for less than a full history.

Information is rarely neutral. Sometimes it is private, dangerous, embarrassing, offensive, trivial, or sensitive in some way. Other times it is important and vital. Knowing when to seek empirically adequate beliefs often requires one to recognise the *value and nature* of the information in question.

Finally, this virtue manifests itself in individual agents who are part of a group or community. Full knowledge—or more true beliefs—may be superfluous or needlessly taxing on an agent’s cognitive resources because of the presence of other competent agents. For example, a man may want a phone number he sees on a billboard as he and his wife drive by it. Yet, he might acquire the first six digits only. He may then ask his wife to acquire the last 4 digits. There is no need for him to know the entire phone number. Instead, he can offload the epistemic

¹¹³ Nosiness includes indiscriminately seeking true beliefs.

burden to another agent.¹¹⁴ In this case, full knowledge is being obtained, but not by the individual. The individual is purposely avoiding the acquisition of relevant information and true beliefs, and doing so virtuously.¹¹⁵

This commonsense virtue is frequently practiced by agents in groups who are trying to solve a particularly complex problem. Suppose, as is often done, we divide the epistemic labour between people in a group so that no agent gains full knowledge of the solution, only parts of it. The group may exhibit full knowledge and understanding, but the individual agents do not; it may not even be efficient or possible for those comprising the group to have full knowledge or understanding. Suppose the group is working on a theory requiring specialists in a diversity of topics. For example, a group of specialists may be working on a metatheory to explain some set of phenomena. Perhaps it requires a biologist, physicist, and chemist, each understanding only her own area. Or take the following examples by way of analogy. A victim of a car accident may require a neurosurgeon, cardiothoracic surgeon, and orthopaedic surgeon. None of these doctors has full knowledge or understanding of the entire surgery; instead, they specialise to complete their particular task.¹¹⁶

¹¹⁴ He is also devoting resources to driving and has limited time to acquire all of the information as he drives past the sign.

¹¹⁵ Notice that a proper evaluation should not be on the discrete beliefs obtained by Jones or his wife. If so, then we will make the wrong evaluation. Anyone with commonsense will recommend that Jones and his wife split the epistemic burden.

¹¹⁶ Even in more simple surgeries the surgeon does not have full knowledge of the entire surgery. The surgeon is aided by nurses and anaesthesiologists.

Likewise, suppose we wanted to build a house. Perhaps no one person understands all aspects of building a house. Construction workers, plumbers and electricians all know different aspects of house building. The electrician on her own has no clue how to build a house, neither does the plumber, nor any one construction worker. Only together do they know how to build the entire house. In fact, we often do not want people with such widespread knowledge. We often choose someone to do the electrical work who specialises in the area. The same is true of surgery. Most people would prefer their brain surgery done by a neurosurgeon as opposed to a general surgeon. In fact, attempting to understand or know too broadly may interfere with one's functioning as a member of a team of specialists. When one is part of a team, it is often necessary to discriminate between the truths one needs to know and others that can be ignored or superficially understood.

In the same way, sometimes full knowledge of some theory, event, or idea is not attributable to any individual, only to a group. Moreover, seeking a full understanding may hurt the individual's performance in her area, which may in turn hurt the group's understanding or capacity to solve a problem. So it may be an epistemic shortcoming to seek full knowledge.

Sometimes knowing when to look for empirically adequate beliefs entails looking for false beliefs. Where true beliefs might be too costly, false ones can be particularly useful. This happens in explanations of some scientific theories, which would be quite difficult to understand if illustrated with sets of true

propositions. For example, one may understand the movements of planets by assuming that each of them is a mass point and that planets do not have gravitational effects on one another.

Take the following two cases in which agents appear to be virtuously seeking false beliefs. Suppose I want to learn to be an expert poker player. But I know that the best poker players are good at detecting and calling bluffs. So I hire an expert to play poker with me. The best way for me to become resistant to bluffs is for me to get bluffed. I do not want some amateur trying to fake me out. I want to buy into the bluff so that I can get better at detecting them in the future. So I purposely go out of my way to set up a situation in which I will acquire false beliefs. Similarly, I may want to obtain a skill in detecting subtle sophistry. So I ask some clever sophist to have a conversation with me, inviting him to mislead me. So we argue about a variety of topics. Sometimes I think I have won the discussion, other times I am not so sure. At the end of the session I exclaim, "I think I have probably come to believe a few of your lies." And I am quite pleased with that result.¹¹⁷

Some environments, especially socially constructed ones, present problems that require only empirically adequate beliefs to solve. This sometimes includes ambiguous, incomplete, and even false beliefs. But they are beliefs that allow the agent to accomplish some intellectual task or solve relevant problems.

¹¹⁷ Adam Morton raised this example.

The point above can be summarised by the following argument. It implies a plurality of value because there are many kinds of problems to solve:

1. Problem solving is an epistemic task.
2. Solving a problem is epistemically praiseworthy/valuable.
3. True belief and error avoidance are not required to solve every problem.
4. So true belief and error avoidance are not required to achieve epistemic praise/value.

Knowing when to Forget

When an agent has particular epistemic aids present, it allows her to employ virtues not otherwise employable. Some aids allow the agent to use her environment to augment her intellectual powers or offload tasks and beliefs. How this is done will depend in part on the aid. The presence of another agent is one kind of epistemic aid. In such cases, it can be virtuous to offload the epistemic burden onto other agents if doing so allows one to preserve resources for other intellectual tasks that could not otherwise be accomplished. The case of the phone number demonstrates not only the virtue of knowing when to seek an empirically adequate belief, but also the virtue of knowing when to forget true ones. The husband may count on his wife to remember the number if his attention is focused on something else.¹¹⁸ Or take the common practice of one

¹¹⁸ A physical analogue to this is asking a friend to help carry a heavy box so that the strength required is divided between two people.

person requesting to be reminded of some event or belief: “Remind me to pick up beer for the party”, “Remind me the next time I order salad to ask for no peppers”. Sometimes the request is made as a sort of “failsafe”, *in case the agent forgets*. However, other times the request is made so that the agent asking for the reminder does not need to think about it or bother retaining the belief. That is, the request is often made *so that the agent may forget*.

The ubiquity of the virtue of forgetting¹¹⁹ has motivated us to develop additional aids and strategies to improve our performance of forgetting. For example, we have invented day planners, journals, diaries, electronic alarm clocks, a string around a finger, and writing on a hand.¹²⁰ Notice that these tools *promote* the act of selective forgetting.

These aids are frequently used in education. Professors write lecture notes and outlines, freeing up resources so that they may forget the particular order and content of information. Students take notes during lectures. Good students discern the information and then proceed to *write it down*. They do not attempt to retain all of the information being presented. Notice that they do not even attempt to retain all of the *important* information. Instead, it is kept external to the agent, put out of her mind, and retrieved later when it is appropriate.¹²¹ Students forget particular true beliefs and store them on paper

¹¹⁹ The virtue of forgetting is often related to the virtue of intellectual trust. Virtuous agents know when to trust the testimony, memory, and cognitive capacities of others such as specialists.

¹²⁰ A physical analogue to this is using a lever to augment one’s strength.

¹²¹ Some information of course will need to be retained, not only for exams but further performance in life.

and computers, which is somewhat similar to offloading these beliefs onto other agents. Moreover, sometimes we seek information—often particular details—that either we know is not worth remembering, but serves a short-term purpose, or we lack sufficient cognitive powers to remember. All of these amount to virtuously forgetting true beliefs. We do this because we recognise there is often little need to remember such a high number of true beliefs, even temporarily relevant ones. This should not bother us. We forget most of what we experience.¹²² Virtuous agents not only understand this is normal for beings with limited intellectual powers, they use this to their advantage in order to accomplish mental tasks by adopting strategies to control which truths are forgotten.

Finally, another reason to forget is that a belief obtained may have been obtained from private or privileged information. Or it may be that some information is particularly distracting or upsetting in such a way that threatens future epistemic endeavours. So the *kind* of information or belief will play a role in determining when it is virtuous to forget.

This reason to forget also justifies avoidance, which is another virtue. It is sometimes virtuous for agents to avoid acquiring true beliefs because the information is of a particular nature. For example, one can virtuously leave the room when particular types of information are being presented.

¹²² It might even bother us to remember some things we do not want to remember.

Knowing when to Ignore Relevant Low-Level Problems¹²³

Another commonsense virtue that is related to both the virtue of looking for empirically adequate beliefs and forgetting is a sort of intentional ignorance. There is much information that agents might be well advised to ignore. One of the more common reasons to ignore information is that it is not particularly relevant. For example, a detective investigating a violent crime may ignore parts of a victim's testimony. Good detectives are selective both in seeking and retaining information. However, sometimes virtuous agents ignore *relevant* information to current problems. Here are a couple of examples in which agents exhibit a commonsense virtue by ignoring relevant low-level information. Suppose a government official, Joan, is tasked with solving the city's traffic problems during peak hours. She analyzes patterns of traffic, which routes are currently available, which alternatives are plausible, and proposes a solution. However, a co-worker, Leroy, objects to the proposal on the grounds that motion is impossible. To Joan's surprise, he has what looks to be good reasons for thinking motion does not exist. He gives a Zeno-like *reductio ad absurdum* argument against the existence of motion and consequently traffic. If Joan seeks to solve the traffic problem, she does not need to, nor should she, entertain

¹²³ There is also a virtue of knowing when to ignore more general relevant information, particularly in cases where one is in over one's head or the information is sensitive, distracting, cognitively paralyzing or destructive in some way.

responses to Zeno's paradoxes. These paradoxes¹²⁴ are certainly relevant to traffic. After all, if Zeno was correct, there is no such thing as traffic since vehicles on his view do not move from one location to another. Joan may be stumped by these subtle and ingenious arguments. However, she would be perfectly virtuous in ignoring Leroy's objection even though it is relevant to her traffic problem. Similarly, one need not solve other low-level philosophical problems such as scepticism and moral dilemmas involving runaway trolleys to form strategies of reliable belief formation and practice medicine respectively.

A more common occurrence of this virtue is displayed by agents who only need a minimal understanding of how a car functions in order to drive. One need not be a mechanic or understand the low-level functions of a vehicle to understand the high-level ones. But notice the low level functions are relevant, even if not useful for completing particular high-level tasks. This sort of ignorance is both common and virtuous. And we would justifiably think an agent is epistemically defective if she always insisted on knowing low-level facts before seeking to know high-level ones.

Scepticism

Before moving on to paradoxical virtues in the next chapter, I want to point out briefly the relation between the virtuous commonsense agent and the

¹²⁴ Theories of time are also relevant.

traditional problem of scepticism. While I do not think scepticism is going to disappear any time soon from epistemology, it can be disarmed to some degree when we consider the view that has formed of the virtuous epistemic agent so far.

I used to think that non-philosophers were intellectually naive in claiming to know such a variety of things. Even claims of perceptual knowledge seemed to me unwarranted. Do they know they are not dreaming or in a computer simulation? How could they? They have not even entertained such scepticism! Some people briefly entertain scepticism in the classroom and then walk out the door without worrying about vanishing into thin air or falling through the floor. We might take a Humean approach and think this is nothing more than unjustified habit. However, if the above intellectual practices are in fact praiseworthy skills that virtuous agents employ, then we can give a more favourable evaluation than Hume's. *Virtuous epistemic agents know when to ignore, forget, and set aside relevant information and low-level problems.* And scepticism is such a low-level problem. Virtuous agents will continue to solve high-level problems, collect high-level evidence and information despite not having answers to low-level ones. It would be a deficiency for an agent to get "hung up" on the low-level issues. Consider the following.

Suppose Fred wants to take his son to the zoo and show him around to experience the variety of animals they have. Fred and his son go to the monkey cage, the horse stables, and petting zoo. Then they come across a new area and

Fred tells his son the striped animals they are looking at are zebras. Now suppose a stranger antagonistically accuses Fred of ignorance, "...something's being a zebra implies that it is not a mule...cleverly disguised by the zoo authorities to look like a zebra. Do you know that these animals are not mules cleverly disguised? If you are tempted to say "Yes" to this question, think a moment about what reasons you have, what evidence you can produce in favor of this claim. The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised to look like zebras."¹²⁵

Suppose the stranger is correct; Fred cannot tell the difference between a zebra and a painted mule. Should Fred refrain from giving his son a tour of the zoo? Is he ill-equipped for such a task? Should he remove himself from his current environment to one that is intellectually safer? Should he refrain from forming beliefs about the animals? Of course not. He does not need to know the animal is not a painted mule in order to guide his son through the zoo. Fred does not need to solve what can be identified as a low-level problem in order to adequately perform, intellectually, on a high-level.

Now, it may be that if Fred were giving a tour to a group of zoologists or philosophers that he *would* need to address the challenge of the stranger; he would need to be able to distinguish between zebras and mules. But the virtue of knowing when to ignore these micro problems, like other virtues, is context

¹²⁵ Dretske, Fred. 1970. 1015-1016.

dependant. Surely he does not even need to entertain such scepticism, let alone have an adequate reply to the stranger, in order to show his son the variety of animals at the zoo.¹²⁶

Likewise, consider the threat posed by the problem of induction to carry out inquiries. If the sceptic points out that inductive conclusions often turn out to be false, that the basis of induction is habit, not reason, and that “experience provides no defensible basis for prediction,”¹²⁷ is it the duty of her audience qua intellectual being to solve this problem or abandon the use of induction before carrying out any further inquiries? I maintain that it is to a person’s credit that her inquiries do not come to a grinding halt when stumped with the relevant problem of induction.

Virtuous agents do not need to meet the challenge of scepticism in order to be epistemic exemplars. In fact, it would be terrible to advise someone to solve epistemic problems “in order”, from the ground up.¹²⁸ That is not to say it is virtuous to ignore the problems entirely. There are potential payoffs to purposely putting oneself in over one’s head. One might stumble upon something unexpectedly. One might have uncharacteristic sparks of ingenuity or make mistakes that reveal strategies for a solution.¹²⁹ Or it may be that, while it is not virtuous for the individual agent to worry about a particular problem, it *is* virtuous—or even necessary—for a community of agents to worry about it.

¹²⁶The tour will include imparting many true beliefs and correcting false ones.

¹²⁷ Hookway, Christopher. 1990. 165.

¹²⁸ Micro to macro.

¹²⁹ It may eliminate a poor strategy.

Excellence and Fitness

Monism assumes the only thing we ought to be concerned about qua intellectual beings is the truth goal. We are situated, embedded, embodied, and bounded people acquiring information that is often not neutral. Our evaluations reflect our awareness of these features.

Other domains of evaluation are likewise indexed. Consider the concept of fitness.¹³⁰ Fitness is attributed to an animal when one wishes to describe its physical performance—or potential performance—in a particular environment. If one were to change the environment, the fitness description may also change. So a fit animal in our world may be unfit on Twin Earth. “Fitness of an animal may involve internal states, but the value of those states is determined by reference to the external environment.”¹³¹ Julia Driver points out that sharp teeth, for example, are indicators of fitness only in some environments.

Fitness is also indexed—to some degree—to the species of the animal. If a species is typically able to hunt small mammals efficiently and run at a particular velocity, then we would withhold an ascription of fitness when some individual animals of that species lacked those capacities. If fitness is a sort of physical excellence that is indexed to both the species and environment, it would be unhelpful to develop a notion of fitness that is for the most part out of reach

¹³⁰ I am not using the term here as an evolutionary biologist might.

¹³¹ Driver, Julia. 2000. 130.

for the species. Tigers do not fly. Horses cannot climb trees. And human beings lack the agility of cats.¹³² If plants are the only things on Twin Earth to eat, then sharp teeth will not contribute to fitness. Evaluating physical excellence—fitness—needs to be grounded.¹³³

Epistemic excellence—or epistemic fitness—also needs to be grounded. Just as it is uninformative to evaluate the physical fitness of an animal in an unnatural environment or against a background of idealised capacities, we would be mistaken to think human excellence—either physical or intellectual—ought to be evaluated from Plato’s heaven. On the contrary, it is to our credit as intelligent beings that we as a species are particularly skilled at understanding our limitations,¹³⁴ understanding our relation to our environments, and manipulating the world to accommodate those limitations.

Many of our epistemic evaluations employ virtues that are not always truth seeking, truth preserving, error avoiding, and error removing. Instead, they aim at cultivating a kind of agent that is flourishing throughout life. Just as physically fit people achieve more than momentary feats of strength and endurance—they have healthy *lives*—so too are virtuous agents developing an enduring character that is more than momentary feats of truth acquisition and

¹³² There are exceptional cases of course. But we do not define fitness by the exceptional cases; we define it by the usual ones.

¹³³ Moral virtues are also set against a background of humanity’s situatedness, embeddedness, embodiment and boundedness. We do not think it cowardly of someone who did not even attempt to stop a runaway train. Instead, we think it foolhardy of someone to attempt such a thing.

¹³⁴ Fitting here is the Greek aphorism “γνώθι σεαυτόν *gnōthi seauton*”, know thyself.

error avoidance. Commonsense virtues contribute to this kind of agent. However, they are not the only ones. Some epistemic virtues are rare or specialised. Some even look like vices. I turn now to those virtues.

Chapter 3

Finite Agency and Long-term Success

The virtue epistemology being developed is one that acknowledges the concrete agent, as opposed to an abstract one. At least three broad observations can be made about real world agents. First, people are finite. There is much we can say about the limits of human capacities, skills, talents, and potential. We have already noted that agents are situated, embedded, bounded, and embodied. We can of course say much more about what each of those entails, which will inform a theory of virtues. Knowing *how* we are limited provides insight to the kinds of virtues for which we can and ought to strive.

Second, we *vary* in our limitations. Take for example the limitation of embodiment. Although we are all embodied, our bodies are not equally limited. Bodies vary from each other, sometimes significantly; there are gender and sexuality differences, for example. Embodied agents encounter the world literally from a particular angle; and these bodies and angles are not the same among agents. They often vary considerably, which will determine what is possible and what is not.

We also differ in cognitive powers. One person may find it easy to accomplish tasks in a given order: x, y, z. Others may do better completing only part of x, before moving on to y or z. The diversity of our limitations often entails

using a variety of strategies.

What these differences mean for the viability of a unified virtue theory goes well beyond the scope of the present work. Nevertheless *how* we differ from each other ought to inform a theory of virtues, especially since it plays a central role in our everyday epistemic evaluations. Good sense for one agent is often bad for another. Some people can take shortcuts and still solve the problem, acquire the belief, or navigate their environment. Others need to take things more slowly and systematically. In short, heuristics that work well for one person may work less for another.

Just as it makes little sense to put forth a theory of virtues for human beings that assumes vast computational powers, it is equally unrealistic to ignore the *variety* of skills, properties, and capacities of people. The consequence of this observation is that the kinds of epistemic characters that are good for self-development vary. Thus, we have goals of being the kinds of effective agents that we in particular can become.

Once we acknowledge these general aspects we can account for an entirely different kind of virtue often found in epistemic evaluations. These are paradoxical virtues. Unlike the standard variety, the traits, procedures, processes, and behaviours employed in paradoxical virtues sound less than noble. Some even sound like well-known vices. That might appear bizarre. Indeed it is. However, it is entailed by facts about the human condition. Below I will present paradoxical virtues that are identifiable as such *only because* we are

the kinds of beings we are. If we had physical and cognitive capacities far beyond our present ones, we would have no need for such virtues. Instead, we would have only “pure virtues”, untainted by the presence of flaws and deficiencies. But we do not have such far-reaching capacities, so we often need “less noble” virtues.

Coupling Constraints

Before identifying some of these virtues, I want to make a few general observations about their paradoxical nature. All virtues are somewhat conditional; that is, they are indexed to particular features of the agent and her environment. Furthermore, virtues are essentially normative; they are good to have. However, as already noted above, virtuous people often adopt less than ideal heuristics when in the grips of particular flaws and deficiencies. And sometimes these strategies and behaviours look like vices, especially if viewed out of context.

Take the following non-epistemic cases in which particular capacities, skills, and strategies are available only in light of existing flaws or conditions. Suppose two hyenas wish to eat a lion cub that is being guarded by a lioness. The plan is simple: one will act as bait while the other one sneaks up and takes the cub, hopefully before the lioness becomes conscious of the deception. Now, the plan will not work if the lioness realises there are multiple hyenas, or if the

baiting hyena does not come across as a threat that is seriously vulnerable. It is not a good thing to be the prey; and generally it is not advantageous to be inferior. However, the inferiority of the hyena, the fact that it can be regarded as vulnerable prey, allows the hyena to be bait. If the hyena did not have those relative flaws, it could not pull off such a deception. Suppose hyenas were more threatening on their own to both the cub and its mother; the lioness might not stray so far from her cub. Instead, she might stay much closer, have less confidence, be more guarded, and be aware of the magnitude of the threat. So while it might be better, generally speaking, for a hyena to be stronger and more threatening, its inferiority allows it to take up a strategy that is sufficient for achieving particular goals.

Consider the following strategy adopted by an injured boxer. Suppose boxer x has a cut above his right eye. There is a real danger, both of blood pouring into his eye and the referee stopping the match. But boxer x is well aware of his disadvantage; and he knows that any further punches he takes to the injured area will surely result in a loss. Knowing this allows him to take up a strategy he could not otherwise take. He purposely drops his right glove just enough to make boxer y think he is being sloppy in guarding the right side of his face. Boxer x knows the injured eye is a high priority target for his opponent; so he is purposely creating an opening to bait his opponent. Attempting to hit the injured eye will expose his opponent's left side, allowing boxer x to get some much needed punches in on his opponent. Injuries are not good in boxing, and

leaving an area unguarded is not an ideal strategy for winning, but they can be used when necessary to achieve goals. *X* is taking a risk of desperation: it might save the match for him, and it might give a rapid loss. But given his present condition, this is his best shot of winning.

It is a sign of intelligence that one makes the best out of less than ideal conditions. The boxer has learned this; and in some way the hyenas have learned to do something similar. Paradoxical epistemic virtues are similar in that they exhibit intellectual skills in unfavourable environments. The following are some candidates.

The Virtue of Dogmatism

We identified in chapter one the virtue of open-mindedness, which includes, among other things, a disposition toward oneself as a limited and flawed inquirer. Open-minded agents regard some of their beliefs to be flexible and open to revision under appropriate conditions. Consequently, some have identified dogmatism and gullibility as the extremes—vices—to which open-mindedness is the mean. However, dogmatism can be virtuous. To be clear, an agent is being dogmatic if she holds some set of beliefs firmly without any consideration whatsoever for the evidence of competing views, which she is aware exists.

Just as there is little hope of creating a list of conditions for which one can

employ more standard virtues, I see little reason to expect such a list for paradoxical ones. However, the cases below are reasonable candidates for cases where dogmatism is being employed virtuously.

Suppose Tom is a scientist at a research institute. He is attempting to come up with a unified theory to explain a variety of phenomena that theorists have found puzzling. Tom's research is based on widely endorsed work done before him. He is making steady progress and is in the final stages of completing a unified theory. Now suppose his colleague Karl is a famous metaphysician and claims he has written a manuscript that must be read by Tom before he does any further research. Tom briefly looks at the table of contents of the manuscript and deduces from it that Karl's work aims to undermine scientific theories and methods that are almost universally accepted. Since Tom's research depends on the veracity of these theories, Karl's manuscript also aims to undermine Tom's research. We can imagine that if Karl has a reputation for academic integrity and cleverness that Tom *may* consider reading the manuscript before completing his research. But given the radical nature of Karl's manuscript, the progress made, the years of research invested, and the imminent completion of a unified theory, Tom may reasonably dismiss the manuscript out of hand.¹³⁵

Suppose Karl objects to Tom's dismissal; he claims the views put forth in his manuscript are revolutionary and challenge the very foundations of science.

¹³⁵ Tom may treat Karl roughly the same way in which he treats sceptics. Hookway discusses the problem of epistemic autonomy in light of the one's capacity to reflect and question the reliability of one's faculties and methods of inquiry. Hookway, Christopher 1990. 146-168.

Furthermore, he claims to have no less than nine arguments, each one independently and sufficiently, according to Karl, proving his radical conclusions. Tom may reasonably hand the manuscript back to Karl, resisting the challenge put forth to become acquainted with any of his arguments, and continue his research.

The critical reader may want to insist that Tom's dogmatism is not intellectually warranted, even if it is pragmatically or morally warranted. However, this is question-begging. On the view being put forth, he may be warranted in all three ways simultaneously. To suggest otherwise is to endorse the view that the intellectual is mutually exclusive with other domains of human life. But why is *that* kind of intelligence worth having at all? If intellectual virtues are valuable, their value must depend on our lives. Once you remove the life,¹³⁶ you remove the value. In the above case, Tom judges that completing his research is a better use of his intellectual capacities than responding to Karl's manuscript.

Virtuous dogmatism is not only common, but sometimes consciously welcomed by people. For example, professors are dogmatic when they dismiss out of hand challenges made by students. This is not uncommon in introductory courses in epistemology. An eager sceptic might challenge every knowledge claim put forth far after the unit on scepticism is completed. At some point, the professor may simply ignore a student's hand because she anticipates a sceptical

¹³⁶ That is, the different domains and environments that make up one's life.

challenge to the content currently being discussed, perhaps the contrast between rationalism and empiricism. We can imagine the student speaks out of turn, insisting the entire discussion is moot given that we do not know anything anyway. Not only is it virtuous for the professor to dogmatically reject the sceptical challenges, but the dogmatism will be welcomed by many students in the class who are hoping to learn. What the professor and other students understand is that not every challenge needs to be addressed before progressing through a discussion, lesson, or solving further problems. This would be virtuous behaviour by the professor if she has failed to convince the student to let the misguided discussion go. So there are conditions under which it is virtuous to be dogmatic.

David Lewis once remarked that dogmatism is sometimes necessary for scientists.¹³⁷ There will always be extremists challenging even the most established scientific theories. Sometimes extreme views give insight into the problems of dominant theories; other times they turn out to be correct. However, there are times in which it is perfectly virtuous for a researcher or inquirer to dismiss challenges that are presented without any consideration whatsoever. The alternative is supposing researchers have the responsibility to address all challenges they are confronted with, even challenges to widely supported and accepted theories. We might wonder if anything would ever get done under such a view of intellectual responsibility. Does the astrophysicist

¹³⁷ Adam Morton mentioned this philosophical anecdote in conversation.

really need to address the conspiracy theories surrounding the moon landing before working on the details of a proposed space station? Should we give flat-earthers—conspiracy theorists who believe the earth is flat and all evidence of its spherical shape is fabricated for mass deception—much or any attention? No. It would be a waste of one's time and epistemic resources to entertain such objections whenever they are raised.

In more ordinary cases, we practice this kind of dogmatism when we refuse to respond to telemarketers who ask us why we cannot take the vacation offer—we know they have a script to address any response we give. When someone is trying to get us to buy a warranty for a recent purchase, we do not actually think we need to respond to every argument they have. We will often dismiss them all out of hand. We might walk by someone who claims she has proof the world is ending in three weeks. We do not think we have any obligation, intellectual or otherwise, to find out what evidence that person possesses and evaluate its strength.

Responding to challenges confers upon them a certain amount of credibility. However, some comments, questions, challenges, and arguments deserve no such credit. Responding to something gives the impression that it was worth responding to, which is often not the case. Thus, dogmatism is exactly the posture a virtuous agent ought to take toward such challenges. This can usually be done politely.

Sometimes we need to be dogmatic just to get stuff done. This may

require taking up a dogmatic position for an extended period of time. Tom, for example, may need to be dogmatic just to see where his research takes him. Perhaps he could not accomplish that if he allows himself to be challenged by Karl. And this dogmatism may last months at a time, perhaps years depending on his research. Now, this does not mean that *everyone* needs to be dogmatic concerning challenges. After all, Tom is not the only competent reader; someone else can read Karl's manuscript. So Tom may be dogmatic knowing that some other person or group can just as easily take interest in Karl's arguments. Tom is at a sensitive stage in his research. So he is more than justified being dogmatic. Finally, notice that the above cases of virtuous dogmatism often have explicitly epistemic aims. In the case of Tom it could be something as traditional as truth seeking. Or perhaps Tom wants to finish his project because the research itself is central to him becoming a particularly adept inquirer. Tom may have little optimism for his project, but perhaps he sees value in acquiring the skills associated with completing it. Research is sometimes done for the training, which may or may not depend on securing true beliefs. Or perhaps this research project is required of him to open up future opportunities to become the kind of inquirer he is hoping to become. In the case of the classroom sceptic, the professor's dogmatism is aimed at refining and cultivating learners. She has a classroom full of inquirers interested in refining their abilities; and dogmatism in this case allows her to accomplish that task.

The Virtue of Self Deception

There is much literature that tries to say carefully what self-deception is. I do not intend to interact with that literature. But I am interested in cases of self-deception that are epistemically praiseworthy.

There are several ways in which an agent may deceive herself. Let us describe one type of self-deception in the following way: one is self-deceived about belief p if one has the dispositional belief that p and attempts—implicitly or explicitly—to stop p from becoming an occurrent belief. Alternatively, one is practicing self-deception just in case one has a strategy for avoiding coming to the belief that p —because one suspects that there is strong evidence that p , so one has to be careful not to meet it. There is often a temporal aspect that is involved in self-deception. One may know x at time t_1 and then seek to believe $\neg x$ at t_2 . To get some idea of the kind of self-deception I have in mind, I will first identify cases I take to be examples of self-deception. I will then propose cases of virtuous self-deception.

People are practicing self-deception when they fail to acknowledge evidence that threatens beliefs they want—or need—to retain. For example, a man may ignore relevant evidence that his spouse is being unfaithful. Perhaps he smells men's cologne on her clothes from time to time. He may have run into her while she was with another man in a bar. He dismisses flirting as nothing but friendliness. He attributes her diminished desire for sex to stress.

A paradigmatic case of self-deception is the alcoholic who thinks his drinking is under control. Often this sort of deception is not only held by the alcoholic, but also his friends and family. It is unpleasant to acknowledge addictions in oneself as well as in loved ones. So it is often not a belief that is easily accepted without some attempt to undermine or trivialise it.

There are five stages of grieving: denial, anger, bargaining, depression, and acceptance. The first stage is often identified as self-deception. For example, when a patient is diagnosed with a fatal disease she may immediately accuse her doctor of incompetence, despite years of confidence in her physician. Or take the case in which a woman has been rushed to the hospital because she is in labour. Only after talking to a doctor does she discover she is pregnant. We may be charitable to some women; perhaps the bump in her belly was not very noticeable. Perhaps she took many of her symptoms to be evidence of stress and illness. But surely some of these women, despite overwhelming evidence, are engaged in systematic self-deception. The motivation behind such denial can vary. And to some degree we might sympathise with the self-deception that some young women practice over several months.

The above cases give a general idea of the phenomenon I have in mind. And they tend to be regarded as cases where someone is in the grip of a vice, a failure to acknowledge a truth for which one has already been given sufficiently strong evidence. However, notice that those typical cases of self-deception deal with truths we think the agent *ought to acknowledge*. The alcoholic will, in the

end, probably benefit more from acknowledging with his addiction. Coming to terms with a life-threatening illness, as unpleasant as it may be, will likely result in better last days. We think the pregnant woman, while perhaps sympathising with her fear, ought to be prepared to deal with her child one way or another. Thus, not only do we identify self-deception as a vice practiced by some people, but I suspect we identify it as a vice because we have already determined the value of the truth being avoided, namely that it is a truth worth believing.

But suppose it was not such a truth. Might we think of cases where the agent is better served to deny epistemically harmful truths? I will now offer a few examples of virtuous self-deception. These are cases in which an agent reasons away, dismisses, reinterprets, or even tampers with strong evidence that threatens to cause undesirable and perhaps epistemically harmful true beliefs.

Victims of tragedies are not expected to come to grips immediately with their experiences, certainly not in every context. Instead, they go through a process in the care of a professional counselor. While the end goal is truth, notice that it is not immediately desirable. In fact, temporary self-deception may be required in order for a victim to understand truths about her past.

Or suppose Donald is an expert problem solver. He has been tasked with solving some problems in Libya. Donald knows from co-workers that Libya is experiencing a civil war, so he fears for his safety. Now Donald knows that if he is too concerned for his safety, he will not be able to complete his task, for his concentration will be disturbed by his worrying. So Donald avoids international

news in order not to expose himself to any reports of violence in Libya.

A common case of self-deception is the manipulation of clocks. A person who tends to be late to meetings might trick herself into thinking it is later than it actually is. My wife does this. She has set each clock in our home and car ahead by four or five minutes. Now she may look at the clock and remember the ruse. But she often does not; sometimes she really does think it is five minutes later than it actually is, which makes her rush out the door to keep her appointments. She has intentionally tampered with strong evidence in order to make herself believe something that is false. She does this to accomplish a variety of tasks.

Suppose I am aware of an epistemic flaw in myself that I find difficult to correct. Suppose x is true and whenever I think x , I infer—for whatever reason—something false.¹³⁸ But I am aware of my tendency to make poor inferences from x . x is not a terribly important truth, so I replace x with a close approximation, x' . x' is literally false. But I make good inferences from x' not x . For example, x may be something like “physicians recommend having eight cups of water per day.” The water contained in food is meant to be included in this calculation. But many people find it quite difficult to judge what eight cups of water look like over the course of the day. So they replace x with x' , “physicians recommend *drinking* eight cups of water per day.”¹³⁹

¹³⁸ Some truths are difficult to use in reasoning. For example, some people find it difficult to compute tasks involving multiples of seven. So they are likely to infer an incorrect answer. They may adjust the numbers to be multiples of 6 if an approximate answer is sufficient.

¹³⁹ It is interesting to note how ubiquitous this substitution is. I suspect most people today are not even aware of the original recommendation. It has been systematically replaced with x' .

Similarly, suppose a scientist seeks to understand some set of events using a theory y . However, y is only true because x is true—quantum field theory. But suppose quantum field theory is too complex and mysterious for her to understand. Fortunately, y is just as coherent if she employs x' —she pictures atoms as miniature solar systems—instead of quantum field theory. She may have serious doubts about x' , and consequently y , if she goes looking for evidence, but she knows something like x' must be true. So she makes no genuine effort to discount x' .

We might think it would be better, epistemically speaking, if the above people did not have to lie to themselves in order to accomplish their tasks. The pluralist can grant this: epistemically speaking, given different intellectual skill sets, the above agents would be better off refraining from self-deception. But given their actual skill sets, it would be *worse*, epistemically speaking, if they refrained from the occasional self-deception.

I want to make an observation about the relationship between virtuous truth avoidance, forgetting, and self-deception. I maintain that, while distinct, there is no epistemically relevant difference between them that would allow one to resist the inference that if one is virtuous so are the others.

Recall that in chapter two we identified truths that agents can virtuously avoid. We acknowledged that some truths have particular qualities that make them unvirtuous to acquire, for example, truths that are epistemically burdening, trivial, private, sensitive, or harmful. It is virtuous, for example, for an agent to

refuse to be exposed to office gossip, information that wastes her epistemic resources, or is harmful to her or her community. So an agent can virtuously avoid obtaining evidence and true beliefs she ought not to acquire. We can express this as the following *a*: Agent (*S*) can virtuously avoid (*VA*) acquiring a true belief (*P*); or *S VA P*.

From this observation we acknowledged that sometimes these sorts of truths are not successfully avoided. So a person can acquire beliefs she should not have. This naturally leads to the virtue of intentional forgetting. We recognise that sometimes it is both commonsense and virtuous for an agent to forget particular truths she should not have obtained in the first place. So from the above observation that *S VA P* we inferred that sometimes *b*: Agent (*S*) can virtuously forget (*VF*) a true belief (*P*); or *S VF P*.

The inference from *a* to *b* is a natural one. If *P* is what we might call epistemically negative, then it does not matter *in principle* that *S* has already obtained the belief. It benefits her to jettison that belief.

However, if it is virtuous for an agent to avoid truths sometimes, and if she fails to avoid them, and furthermore occasionally fails to forget them, then I see no reason to resist the further claim that it is also sometimes virtuous to practice self-deception about *P*. Some truths are hard to forget. So an agent may need to exercise some form of distraction, distortion, and restraint on her beliefs or exposure to evidence. Consequently, *c*: Agent (*S*) can practice virtuous self-deception (*VS-D*) of true belief (*P*); or *S VS-D P*.

So instead of avoiding or forgetting a truth, the agent practices self-deceit with the same goals in mind. I do not see any epistemically relevant difference between avoiding, forgetting, and practicing self-deception about a belief that should not have been obtained in the first place. So virtue of truth-avoidance is related to virtues of forgetting and self-deception.¹⁴⁰

We might note one epistemic difference between truth-avoidance, forgetting, and self-deception: the agent's disposition toward *P*. The agent does not believe *P* in truth-avoidance, but she does when she tries to forget and deceive herself.¹⁴¹ So truth-avoidance involves belief formation while the other two consist of belief revision. However, it is not clear why this distinction is relevant. I should be clear that I am not maintaining that *b* follows necessarily from *a*; nor if *b*, therefore *c*. It may be that in token cases it is appropriate to avoid some truths, but not forget them once obtained. And it may turn out to be virtuous to forget some beliefs, but not practice self-deception concerning those same truths. That will depend on the context.

Nonetheless, once we acknowledge that some truths have negative epistemic worth, we open up the possibility that token truths can be avoided, restrained, or removed from an agent's library of beliefs. Making this observation, we can expand our candidates of virtuous self-deception by revising cases of virtuous truth-avoidance and forgetting.

¹⁴⁰ There is some relation here with the virtue of dogmatism since the evidence being dismissed will sometimes lead to truths.

¹⁴¹ As is the case of the manipulated clock.

The Virtue of Intellectual Sloppiness

Typically, we do not think it is good to be convinced by bad arguments and insufficient evidence. Much of what philosophers do is pick out fallacies and errors in reasoning, including misleading evidence. And virtuous agents are careful about drawing conclusions that are unwarranted or supported by inadequate evidence. However, sometimes the virtuous agent settles for poor and insufficient evidence, comes to hasty conclusions, or develops clumsy theories. For example, sometimes people will strategically adopt folk theories knowing they are unwarranted or largely erroneous to accomplish their particular epistemic ends.

The reason for such a strategy should now be familiar. The ideal strategy—if there is one—is not always feasible. We do not always have the time or resources to develop scientifically respectable theories; and sometimes the only evidence available is poor, unreliable, and insufficient. Even when one already has a good theory or evidence available, one might opt for an incomplete, unwarranted, or folk view that is sloppy by comparison.

Take for example an agent who adopts a folk theory of social behaviour. She makes quick judgements about people's personalities based on their behaviour and facial expressions. The woman with the stern face is unfriendly. The man laughing loudly has low self-esteem. And the teenager wearing too

much makeup is promiscuous. She knows full well that she has very poor evidence to make such inferences. And she would never seriously offer her folk theory up as a competing theory for experts to consider. She may even genuinely argue against such a theory in an academic setting. Reliable evidence may suggest that people are much more complex than her outlook allows. Some people's dispositions appear sterner than others, though they are perfectly friendly. Others do not know how loud their laugh is; and still others do not know how much makeup is excessive. But she may find it useful in a particular social setting to disregard the evidence and the more accurate theories she knows of and adopt a crude view of people, their behaviours, and personalities. So she reads people, attributing properties to them that are often false, at the very least unreliable. While the theory she is applying is sloppy and flawed, it captures enough correct information to guide her. She may need social judgements, for example, to work in an office and finds herself unable to regulate her work-social life with anything more nuanced.

Her success may be social or moral. However, it may also be epistemic; her folk views are deeply related to many judgements and inferences that she makes, which are central to information gathering and belief formation, revision, and maintenance. And while a more complex and supported theory will provide more accurate explanations of her *experiences*, her aim is something else; and she may not be able to spare the resources required to obtain, maintain, and apply such a theory. The folk view, as sloppy and unscientific as it is, is sufficient

to accomplish her other goals. A more nuanced view may be less epistemically useful.

Suppose, for example, that she is an astronomer. She has to assess the reliability of a student's report of a new neutron-star discovery announced on the internet. Should she take it seriously and take the time to check it out? She may decide based on how the student is dressed.¹⁴²

Suppose someone uses her intuition "to read people" and is quite good at it. This person has no formal theory and cannot articulate any systematic way of judging people. According to her, she just uses her intuition. Her intuitions may be riddled with error; and if she spent the time to test her intuitions, do research, and take up a well-tested and refined view of social interactions she might have an improved view and understanding of social environments. But it is not clear she would do any better with such knowledge. She might even do worse if she could not figure out how to *apply* the information. Now perhaps *some* people can; but not everyone should be aiming for the same degree of careful thinking and nuanced theories. Some are better off with crude descriptions and conjectures.

Not every context requires careful thinking. Using such a standard universally is analogous to locksmiths using dynamite to open simple locks or engineers cutting their toenails with microtomes. There is simply no reason to use such powerful tools to accomplish the desired tasks.

¹⁴² This example is from Adam Morton.

In communities we acknowledge that splitting the epistemic burden is our best way to achieve collective ends. It is often good enough that we become only vaguely acquainted with the details of every area—perhaps filling in gaps of knowledge with crude descriptions—and save our intellectual resources for areas in which we wish to specialise.

Virtuous agents recognise that they do not need to be experts in every field to accomplish tasks in the world. Sometimes a sloppy tool is sufficient. Occasionally sloppy tools are better. Sometimes respectable theories are impossible to practice—because of some facts about the agent’s intellectual powers. Take for example a tenseless theory of time in which past, present, and future slices of time are regarded as equally real. According to this theory the experience of time’s passage is an illusion. Under such a view it is erroneous to think that my reaching for a door and applying pressure changed it from being closed to being open. Let us assume that a tenseless theory of time is the only scientifically respectable one. Nevertheless, a student of physics may use a sloppy tensed theory of time, navigating her environment from moment to moment because adopting a tenseless view of her actions is either unnecessary or not conducive to success. Or perhaps she uses the sloppy theory because she finds it hard to understand a tenseless theory of time in the first place, let alone apply it in her life.¹⁴³ She needs to apply *some* view. It will do her little good to

¹⁴³ We could make the theory even more out of her intellectual reach by describing it as a quantum theory of time that invokes the many-world hypothesis. Many people find such

insist she must apply only the correct theory when she cannot even comprehend it.

The pluralist can acknowledge that it would be better, epistemically speaking, if she had the intellectual powers to comprehend the theory and apply it to accomplish tasks. But given that she has no such intellectual powers, the pluralist maintains that it is better, epistemically speaking, if she uses a sloppy, crude, or even false theory that is sufficient to accomplish other tasks, which will no doubt include epistemic tasks. It is not clear a monist can endorse such a strategy since it forgoes the truth goal, not necessarily to acquire different truths.

The Virtue of Dishonesty

Finally, let us consider what might be classified as a collective or group paradoxical virtue—a virtue that exists as a property of the group or relations between its members. There are times in which an agent may be virtuously dishonest about research, evidence, and data, including the work of other researchers. Ironically, many cases of dishonesty are identifiable as virtuous precisely because they are aimed at the truth goal. Suppose for example Smith is a researcher who has gone to great lengths to discover a new theory with more explanatory power than competing ones. The plausibility of his theory rests on

interpretations of quantum mechanics too difficult to comprehend, let alone to apply them in decision making processes.

both empirical data and abstract reasoning. Now suppose there is a gap in his reasoning for his conclusions that is detectable only to the most informed experts in his field. Smith suspects that if he submits a paper on his research the review committee will not detect the gap, even though, once printed, a future reader will likely notice it. However, he is confident that, once recognised, someone will attempt to fill in the gap. It is likely that the gap will not be filled unless many people are exposed to his theory and research; so Smith knowingly submits his paper under the pretext that his evidence is complete.

Or suppose a review board at a conference has picked one of several submitted documentary films to show at an evening workshop for its delegates. The event's planner, Mary, who is a leading expert on the topic of the documentary, views it out of curiosity. Mary notices that the documentary contains clever and subtle elements of distorted evidence and editing such that it misleads all but the most informed audience. Mary reasons that if this documentary is shown to the delegates it will encourage false beliefs and promote wasted research by amateurs. However, Mary cannot adequately expose the errors that make up some of the key points in the documentary to the review board, who do not have the background to understand the subtle manipulations being made by the director of the film. So Mary intentionally misrepresents a section of the documentary she knows will change the minds of the review board. Mary's choice is between being dishonest and letting the film be shown to delegates, which will encourage wasted research and the

acquisition of false beliefs. More formally, suppose an audience believes the following:

1. If x then y
2. If y then x

But suppose they are wrong about the relationship between x and y . The truth values of x and y are mutually exclusive— x if and only if not y . Suppose further that it takes significantly higher intellectual powers than the audience has to understand the relationship between x and y such that the audience cannot reasonably be expected to refrain from making the above inferences. Let us grant that the content of x is trivial while y is important. In addition, an agent knows that if she reveals to her audience that $\neg x$, they will mistakenly infer $\neg y$ —mistakenly because the second premise is false; but it does follow from $\neg x$ and the second premise—and once $\neg x$ is revealed she cannot stop the inference from taking place. An agent may then simply allow her audience to believe x for the sake of getting them to believe y if she understands that they will not believe y unless they first accept x . She may even *argue for* x dishonestly if she feels belief in y is being threatened. That may include concealing evidence or data that implies $\neg x$.¹⁴⁴

The kind of dishonesty in the above cases is sometimes employed and accepted by virtuous agents. We recognise that truth is sometimes more

¹⁴⁴ Notice the similarity between dishonesty and self-deception. The only difference is the number of people involved. This implies a further virtue, that of letting oneself be deceived. The poker example might exhibit this virtue.

misleading or less informative given the environment in which they are expressed or revealed. For example, I have been asked many times if I drink coffee. I almost always lie. I claim not to even though I do. The reason for this is that the truth is more misleading than the lie. If I say I drink coffee, I know that my audience will think I drink it daily, which is both typical of coffee drinkers and false of me. If I deny drinking coffee, my audience will often have a more accurate impression. I will drink it; I like drinking it; but it is not something I drink regularly. If I were a guest in someone's home and asked the same question, I would claim to drink coffee, understanding that the host wants to know if I want a cup of coffee. However, it is the context of the question or inquiry that sometimes demands less than truthful answers. It may require concealing some aspects of the truth. The same story can be told with regard to questions of alcohol. I am often asked if I drink alcohol. I almost always lie because in particular contexts, if I say yes, the inference is immediately made that I frequently get drunk. But this is not true, and I would rather communicate that I do not, as opposed to the more trivial truth that I enjoy a glass of wine from time to time.

While the traditional aims of truth seeking and error avoidance play a central role in virtuous dishonesty they need not be the only goals. For example, Mary, the event planner, was trying to stop her delegates from forming false beliefs and consequently wasting resources on the research it might encourage. However, we can easily alter the case and grant that the documentary does not

contain any false or misleading content but will still lead to wasting intellectual resources in research because the added knowledge that would result from such efforts would be mostly superfluous—further research is of no benefit. Or perhaps another area of research is more important and collective effort ought to be spent elsewhere; but the audience lacks the insight to appreciate this and prioritise their talents appropriately. Broadly speaking, while the documentary might be a perfectly accurate account of its subject, it may nonetheless have epistemically negative consequences.

After my wife saw a documentary on sharks, she came to the conclusion that she never wanted to view another one because she knows that any more documentaries will lead her to fear swimming, even in bodies of water that have no sharks. Filtering information is common. Perfectly accurate information may promote false beliefs, prejudices, and irrational fears. Or it may promote accurate beliefs and rational fears, but hinder one's ability to complete computations, solve problems, and navigate environments.

The virtue of dishonesty is possible because of our psychology. Truth does not always beget truth. Nor does it always improve our capacity to solve problems. We are clever enough to know this. Once we acknowledge the diversity of roles that true beliefs play in our system of beliefs and performances, it follows that sometimes we ought to ignore, forget, conceal, and misrepresent them. Insofar as corporate epistemic goals exist, virtuous dishonesty may be appropriate.

We know full well that token truths are not always worth knowing. And we often safeguard others from forming beliefs about them when we think they lack value or diminish our capacity to be the kinds of agents we want.

The Paradoxical Nature

Why ascribe a paradoxical nature to some virtues? Why, for example, is the virtue of dishonesty not simply the virtue of honesty if honesty is thought to be the mean between, on the one hand, being rude or boastful, and on the other hand being deceitful or dishonest? If virtues are contextual, we should expect that the mean will sometimes look like dishonesty just as in some extreme cases courage requires large amounts of aggression or immediate submission to fear. So why identify these as paradoxical instead of just standard virtues in extreme environments?

One reason is that some virtues are restricted conceptually. We cannot, for example, say someone is being honest when she is intentionally deceiving. Despite being indexed, honesty requires an absence of intentional deception. So identifying the virtue as “dishonest” and paradoxical better captures what the agent is displaying.

Moreover, “standard virtues” have been defined against the background of an *ideally limited* agent. Many philosophers claim to recognise the futility of assuming human beings are perfectly rational with unlimited powers of

inference. Unfortunately, the limitations described or assumed in their theories of epistemic exemplars are still of the ideal sort. Many virtue theories often fail to recognise the degree to which agents and their environments are limited and unfavourable. Consequently, they leave out virtues that agents often need to employ when particular resources are not available. Compare two types of limited agents. First, an ideally limited agent recognises her limitations and accounts for them. The way she approaches problems and inquiries maximises efficiency. While limited, this agent retains a sort of ideal nature because her reactions to her limitations are ideal. She is never caught off guard. Her environment is stable, predictable, and there is a most efficient strategy for problem solving. For example, though she has a particular learning style that does not allow her to study for an exam in a maximally ideal fashion, she is able to tailor her study session very well for her limited learning style. She might add pictures to pages of notes to help cue her memory. And she will have all the time in the world to find or draw these pictures. This sort of limited agent appears to be what many epistemologists have in mind when they describe epistemic exemplars.

But suppose we consider another kind of limited agent. She is often able to recognise her limitations and adjust her heuristics to *sufficiently* solve problems she encounters. Sometimes problems arise that she does not—could not—have anticipated. In such cases, her resources may become greatly diminished so that she must take up otherwise inappropriate strategies for

successfully navigating her environment. For example, she may find herself needing to study for an exam with very little time. Had she more time she would begin at the start of her class notes and review them, perhaps rewriting them in order of date or concept or whatever best suited her learning style. Given that she finds herself in an environment with little time, she skims her notes from the start to the middle, then back to the start again, without reading her most recent notes since they are the most fresh in her head. Or perhaps an unexpected blackout nullifies one's "standard capacities" to navigate the room and so one must adjust the best one can. My own view is that our common observations of epistemic exemplars resemble the second agent described.

It is not enough that philosophers recognise the limitation of agents; they must also appreciate that much of one's environment is beyond anyone's control. So sometimes a person's talents and skills appear paradoxical. An ideally limited agent in a predictable environment probably would never need to be dishonest or conceal the truth. She would not need to be sloppy in her thinking, lie to herself, or refrain from exposing herself to evidence, information, and experiences. So relatively speaking, the above are paradoxical; they *sound* like they are vices and epistemically irresponsible; and they would be for ideal agents—even limited ideal ones.

Paradoxical virtues must be identified against the background of the agent's telos, which is sometimes restricted, changed, or even eliminated by her environment. She may sacrifice some immediate truths for distant ones. Or she

may even put herself in an environment that is likely to entail errors for the sake of learning to become an efficient or sufficient inquirer. Traditional theories tend to assume truth is always available or worth obtaining; but that is false.

Much of what shapes and informs an agent's telos is the way in which she must *manage* her limitations.¹⁴⁵ Intelligent people manage and compensate for their limitations by adjusting and adopting goals that are feasible. Some consequences of applying virtues of limitation management can seem paradoxical, for example, no longer aiming at truth or honesty.

In terms of limitation-management, we can say that many virtues seem paradoxical because they have been defined traditionally by their relation to the truth goal. It is not surprising that virtues aimed at non-traditional values appear paradoxical or wrongheaded. So dishonesty sounds like the kind of thing we ought to avoid. However, the degree to which we embrace value pluralism and redescribe agents as situated, bounded, embedded, and embodied, who are often trying to manage their limitations, diminishes the paradoxicality of the virtues.

These conclusions fit well with the general view being put forth. All of the above virtues sometimes sacrifice token true beliefs because doing so figures into the wider telos of the agent. Much of that consists of becoming a particular type of inquirer and problem solver, one capable of efficiently—sometimes only sufficiently—accomplishing both short and long term tasks. These virtues allow

¹⁴⁵ Limitation management is discussed by Adam Morton.

the agent to navigate her environment with less turbulence as it were, which includes, for example, managing the limitations of the self and others in sometimes unpredictable environments. Sometimes particular truths cause epistemic turbulence. So they are intentionally avoided, dismissed, or rejected.

This is not merely *acting* dishonestly, deceptively, and so on, for the occasional payoff. We may recommend that a person cultivate some of the above skills as *traits of character*. A mentor's advice to her apprentice might be, "You'll never understand all these ideas in engineering if you insist on seeing them in terms of basic physics. Instead, persuade yourself that the atom is a solar system, and light is vibrations in the aether, and you'll be a much better engineer. Your success depends on you being able to believe some half-truths."¹⁴⁶ Some of the above paradoxical skills might require forming a type of *character*, as opposed to momentary acts.

Paradoxical virtues are not something that ideal agents would possess or desire; but we are not ideal. That is why they exist in the first place. It is epistemically praiseworthy to take advantage of or compensate for one's limitations and deficiencies, which paradoxical virtues do. And I think this is a commonsense view reflected in our actual epistemic evaluations.

Consequently, it behooves epistemologists to acknowledge that virtues are not always pure. They are context-dependant and agent-relative. Real people are flawed, biased, and constrained in such ways that limitation management is

¹⁴⁶ This example is taken from Adam Morton.

'both necessary and praiseworthy. We are idiots compared to the ideal agent. Fortunately, it turns out the human species as a whole is quite good at being intelligent idiots.¹⁴⁷

¹⁴⁷ Morton has mentioned this idea of being intelligent stupid creatures in conversation.

Chapter 4

So far I have argued that value pluralism gives us the best account of epistemically virtuous performances of limited human beings. Pluralism is already taken for granted and reflected in our everyday evaluations, which often include praise and critique of people's skill in managing their cognitive limitations. Occasionally both commonsense and paradoxical heuristics, which sometimes aim at more than the truth goal, are appropriate. Instead, virtuous agents aim to become particular types of inquirers, which cannot be accounted for by appealing only to one value. I want to continue to build a case for recognising a plurality of value by considering several issues that have become increasingly more important in the literature. Addressing these issues in a framework of virtues suggests value pluralism.

The Value Problems

One of the oldest problems in epistemology—found in Plato's *Meno*—is: *why is knowledge more valuable than mere true belief?* Many epistemologists think that knowledge is more valuable than merely true belief.¹⁴⁸ But how do we account for this added value? There have been many attempts to answer this question since Plato, though all of them are controversial. Some have claimed

¹⁴⁸ Some epistemologists deny the claim. For example, some have argued that knowledge is nothing more than true belief. Others claim that knowledge is valuable inasmuch as it is a form of true belief.

the added value of knowledge is something extrinsic to it. The extrinsic candidate is usually some sort of practical or useful advantage that knowledge is supposed to enjoy. While some have taken up this pragmatic approach to the added value of knowledge, many philosophers, including Socrates, find it unconvincing since true belief appears just as useful in getting things done in the world. For example, a tourist will get to the Statue of Liberty just as easily having true beliefs about how to get there as she would if she knew the directions. Knowing the directions to the Statue of Liberty does not provide any practical value that true believing lacks.

Some have attempted to explain the value of knowledge in terms of its subparts. But this approach carries an added problem; one needs to know what the subparts of knowledge are. No easy task. In the aftermath of Edmund Gettier's attack on the justified true belief account of knowledge, those wishing to address the value problem have adjusted the ancient question to the following: *Why is knowledge more valuable than mere justified true belief?*¹⁴⁹ Consequently, it looks as if one must solve the Gettier problem before one has hope of addressing Plato's, which has become known as the value problem. It is also sometimes referred to as the Meno problem.

It is not surprising that this has become a central concern for value-driven epistemology since the very nature of the problem concerns epistemic value. I

¹⁴⁹ I will use this version of the problem. The reader should feel free to be a minimalist about justification if she feels so inclined.

cannot hope to cover all of the attempts to answer the value problem. However, I want to raise a concern for any attempt to answer the value problem by appealing only to its subparts since it is particularly relevant to pluralism.

Some philosophers are motivated to explain the value of knowledge by its subparts because, regardless of the details of the sufficient conditions for knowledge, true belief is a necessary condition, and true belief is widely acknowledged to be epistemically valuable. The strategy then for many philosophers has been to explain the value of knowledge in terms of its relation to truth. For example, a reliabilist might maintain the added value of knowledge is its reliability. True beliefs can be unreliable whereas knowledge by definition¹⁵⁰ is a reliably obtained true belief. So the added value on this account is the reliability that is provided. And we are to understand the value of reliability by virtue of its truth-conduciveness.

There are a couple of problems with this particular response to the value problem. First, it is dubious that reliable true beliefs are sufficient for knowledge. Consider Russell's clock-watcher. He glances at the clock and sees that it reads eleven o'clock with a sign next to it claiming "certified accurate" with yesterday's date. So he thinks it is eleven in the morning. But the clock has stopped working for weeks and the sign is a prank. His method of telling time is generally reliable and his belief happens to be true. But it is not clear he has knowledge of the

¹⁵⁰ According to reliabilists.

time.¹⁵¹ If the reliabilist thinks there is a difference between general reliability and token reliability, where Russell's clock-watcher is achieving the former but not the latter, then she needs both to distinguish between the two and explain why token reliable true beliefs are more valuable than type reliable true beliefs. If not, then the reliabilist has not identified all of the subparts of knowledge.

This concern can be raised against any attempt to account for the value of knowledge by appealing to only its subparts. Some form of justification, warrant, coherence, proper function, and reliability, among others, are typically identified but regarded as controversial or rejected as sufficient conditions of knowledge. And if they are not sufficient, then they likely will not explain what is so valuable about *knowledge*.¹⁵² That is not to say one cannot solve the Gettier problem.¹⁵³ However, whatever subparts of knowledge are identified—*x*, *y*, and *z*—any such account of value would be as controversial and vulnerable as any *x*, *y*, and *z* theory of knowledge.

One might attempt to bypass the problem by taking knowledge to be a primitive concept.¹⁵⁴ On this approach, instead of explaining knowledge by its subparts, we use it to explain other concepts such as justification, evidence, and belief. Notice on this approach, however, that knowledge remains factive. If I know that *p*, then *p*. But then the problem about the value of knowledge remains. For many *p* it would still be the case that if someone lacks evidence for

¹⁵¹ Russell, Bertrand. 1948.

¹⁵² This objection is raised in Kvanvig, Jonathan. 2003.

¹⁵³ Zagzebski discusses the inescapability of Gettier-type problems. Zagzebski, Linda. 1994

¹⁵⁴ See Williamson, Timothy. 2000.

p , then they do not know that p . So we can still question how the value of knowledge is related to the value of other concepts such as justification and belief.¹⁵⁵

Suppose we ignore for a moment problems finding an adequate account of knowledge. Will that help this strategy in solving the value problem? It appears not, because this strategy must face what has become known as the swamping problem. An example will help explain the problem. Linda Zagzebski has argued convincingly that reliability does not account for the added value of knowledge over true belief because reliable true beliefs are no more likely to be true than beliefs that are merely true. The following analogy makes her point. Compare two cups of espresso made from two different machines. The first cup is made from a machine that reliably produces good espresso. The second cup is made from a machine that has a poor record of making good espresso; it is quite unreliable. Yet suppose both cups of espresso are "good". Is the first cup a better cup because it was made by a reliable machine?¹⁵⁶ It appears not. Likewise, the reliability of the processes involved in achieving a true belief does not make the true belief any more likely to be true than a true belief that was formed unreliably. After all, the merely true belief is already true; it cannot be any more likely to be true than it already is.

¹⁵⁵ Adam Morton pointed out this example to me.

¹⁵⁶ Zagzebski, Linda. 2001a.

Zagzebski's criticism implies a problem for more than just reliabilists hoping to solve the value problem. Reliability—or justification, warrant, coherence—is supposed to add value to a true belief because of its relation to truth: reliably formed beliefs are more likely to be true. Let knowledge, k , be composed of true belief, $T(b)$, and some set of conditions, x , ($k = T(b) + x$). k is supposed to have more value than $T(b)$ because beliefs with properties x are more likely to be true. But notice the redundancy of value being identified. If we add some feature to a belief that is valuable only because beliefs with this feature are more likely to be true, then this is not something valuable *in addition* to the value already present in a *true* belief. The value of reliability, like other features that are thought to constitute knowledge, is parasitic on the value of truth.¹⁵⁷ And since a true belief is already obtained, whatever added value is proposed will be swamped by the value of the true belief. And this is simply due to the fact that a true belief is already true; it cannot be any more likely to be true regardless of the property we add to it, reliability or anything else.

One reason the swamping problem exists is because epistemologists have focused on trying to account for the added value of knowledge that accrues to *the belief*. But once we acknowledge beliefs are not the only objects of evaluation, we can begin to address the swamping problem. Since virtue epistemology conceptually prioritises agents over their beliefs, it is well-equipped to help.

¹⁵⁷ Kvanvig, Jonathan. 2003.

According to some virtue reliabilists like John Greco, Wayne Riggs, and Ernest Sosa, the added value of knowledge over true belief is that the agent who employs her reliable faculties—intellectual virtues—deserves credit. There is something valuable not only in having true beliefs, but in the *grasping* of the true beliefs *by one's own virtues*. The added value on this account is supposed to be accrued to the agent not the belief:

*"We want rather to attain truth by our own performance, which seems a reflectively defensible desire for a good preferable not just extrinsically but intrinsically. What we prefer is the deed of true believing, where not only the believing but also its truth is attributable to the agent as his or her own doing."*¹⁵⁸

The added value of knowledge—virtuously acquired true belief—is accrued to the agent, not the belief. So it looks like this avoids the swamping problem.

Two things are worth noting about this response to the value problem. First, virtue epistemologists of all varieties will need to address the Gettier problem. Insofar as a case can be made that virtuous true belief is not sufficient for knowledge the above account will not have shown what is better about knowledge. Unfortunately for virtue theorists, the Gettier problem is not avoided by introducing virtues. We can imagine cases in which an agent virtuously acquires true beliefs using her reliable faculties, skills, or character traits and yet does not have knowledge. Consider Goldman's fake barn example. Suppose a

¹⁵⁸ Sosa, Ernest. 2003a. 175.

person is driving through the countryside where there are many fake barns and one real barn. An agent driving by would have no clue she is looking at mostly fake barns. She points to a barn and says, "That's a barn." She happens to be pointing to the one real barn in the area.¹⁵⁹ Most epistemologists agree this person does not know she is pointing at a barn. Yet her belief was formed through intellectually virtuous faculties and perhaps character traits.¹⁶⁰

We might even take Gettier's original cases¹⁶¹ to be paradigmatic of virtuous belief formation. The beliefs formed in the original cases were due to impeccable use of logic and reasoning from the best available evidence. There is no reason to think that Smith in the Gettier cases has any defect either in his character or intellectual capacities. He appears to be blameless with regard to his use of deduction in both of the original cases. So there is little reason to withhold virtuous belief acquisition from him. Yet he does not appear to have knowledge in either case. If so, then virtuously acquired true beliefs will not be able to account for the added value of knowledge since knowledge is not virtuously obtained true belief any more than it is justified true belief.¹⁶²

Second, once one attempts to account for the value of knowledge over true belief without appealing to truth-conduciveness, value monism is in trouble. If the value of knowledge is not reducible to the value of truth then there must be more values. Zagzebski seems to have noticed this consequence. She argues

¹⁵⁹ Goldman, Alvin. 1976. 771-791.

¹⁶⁰ Kvanvig, Jonathan. 2003. 84.

¹⁶¹ Gettier, Edmund. 1963. 121-123.

¹⁶² It also seems that someone can get knowledge when acting out of character.

that the rationality of preferring to acquire truth by one's own abilities derives from the contribution it makes to human flourishing.¹⁶³ Her suggestion for virtue epistemologists like Sosa is to abandon value monism.¹⁶⁴ Sosa agrees. Furthermore, in his reply to Zagzebski he denies ever having endorsed value monism. He even goes so far as to argue that there is nothing valuable whatsoever about truth *as such*.¹⁶⁵

I must say, I find the entire exchange between Sosa and Zagzebski rather ironic. For all their talk here of abandoning monism, their accounts of intellectual virtues¹⁶⁶ preclude that possibility. Unless both intend their virtue theories to be confined to theories of knowledge, their accounts are suspiciously myopic. This concern is echoed in Kornblith's criticism of Zagzebski's celebrated *Virtues of the Mind* when he urges a virtue approach to epistemology to use its broad conceptual tools to venture outside traditional topics to investigate other areas of interest.

"Now, to my mind, the prospect of an epistemology modeled on virtue ethics seems exciting in large part because it might change our conception of what the important epistemic projects are and revise our understanding of which epistemic notions ought to be at the center of our concern. For example, a Kantian ethic sees questions about moral principles and duties as central to ethical theory; the ideas of respect for persons and impartiality require analysis because their theoretical importance is at the heart of the Kantian project. Aristotle's virtue ethic does not merely offer a different analysis of these notions. Rather, the

¹⁶³Zagzebski comes close to expressing the view that I endorse regarding the value problem.

¹⁶⁴Zagzebski, Linda. 2004. 192, 197.

¹⁶⁵Sosa, Ernest. 2004. 320. See also Sosa, Ernest. 2003a.

¹⁶⁶That is, virtues for both of them are defined by truth-conduciveness.

*entire Aristotelian moral project is organized around a different question—how to live one's life. . .*¹⁶⁷

Nonetheless, their discussion is fruitful. It implies that value pluralism is attractive or perhaps even inevitable in light of the value problem.

Wayne Riggs has summarised the above tension as follows. There appear to be three assumptions in epistemology that are jointly inconsistent.

1. Knowledge derives all of its value from the ends or goods of cognition.
2. There are only two of these goods: having true beliefs and avoiding falsehoods.
3. Knowledge is always more valuable than mere true belief.¹⁶⁸

The first two, Riggs argues, entail a denial of the third. The reasons for this we have seen above. The value of knowledge cannot be more valuable than true belief if having true beliefs and avoiding error are the only values—and the value of knowledge is reducible to these. So Riggs rightfully claims that at least one of the above is false. His candidate is the second assumption. And from this he attempts to identify a value that is not reducible to the value of truth.¹⁶⁹

I think Riggs is correct that the second assumption is false. But we should not think that because we can identify one as being false that the other two are

¹⁶⁷ Kornblith, Hilary. 2000. 200-201. See also Hookway, Christopher. 2003.

¹⁶⁸ Riggs, Wayne. 2002.

¹⁶⁹ Riggs argues that epistemic responsibility is valuable apart from the value of truth.

true. In fact, I think it can be shown that the third assumption is also false. I will argue for the following claim:

1. Knowledge is not always more valuable than mere true belief.¹⁷⁰

Take for example the true belief that there are n number of motes of dust on my desk.¹⁷¹ Now whatever value this may have, what added value is to be gained from knowing n ? Surely it is not that one has come to believe n because of one's own intellectual capacities. If anything, we are likely to fault someone for wasting her resources trying to acquire such a belief. The value of having this true belief appears to be exactly the same as knowing it: nil. A similar evaluation would be made of people who seek to confirm gossip and rumours.

Now, one may be tempted to reply to these counterexamples by claiming that a necessary condition for knowledge to be more valuable than true belief is that the true belief in question must have some value to begin with; it cannot be a trivial one. So perhaps one ought to understand the third assumption in an alternate way: *Knowledge is always more valuable than non-trivial true beliefs.* However, it is not clear to me that that new assumption is intelligible on a monistic account. The assumption is tantamount to the following: *Knowledge is always more valuable than true beliefs that have some value to begin with.* By virtue of what is a monist able to distinguish between valuable and non-valuable true beliefs, unless she is appealing to some value other than truth? If she is

¹⁷⁰ It will also be shown that sometimes *no belief* is better than knowledge.

¹⁷¹ Sosa, Ernest. 2003a.

invoking some other value, she is no longer a monist. So this alternate assumption is available only to a pluralist. Monists appear to be committed to the view that all true beliefs are valuable,¹⁷² which means she cannot object on these grounds.

In addition, consider the goal of inquiry. It is often truth—not knowledge—that we aim for in our inquiries. We frequently stop our inquiries once we are convinced of some belief. Sometimes we continue inquiring when we see a particular value in also knowing it. But, our curiosity is often satisfied by belief acquisition alone. Maybe knowledge would be more valuable, but that is determined by context.¹⁷³

Sometimes knowledge is actually *less valuable* than true belief. This is the case for at least two reasons. First, the payoff of acquiring a true belief is sometimes higher than knowledge. Second,¹⁷⁴ the cost of obtaining knowledge is often much higher than the cost of acquiring true belief. A mistake many epistemologists have made is assuming knowledge always has a *superior* relation to the agent than mere true belief.¹⁷⁴

Suppose knowing that *p* causes Jones—due to some particular features of Jones' psychology or personality—to become overconfident and sloppy about inferences from *p*. Jones might not make careless inferences from *p* if he merely

¹⁷² Kvanvig defends this view.

¹⁷³ Kvanvig argues that goal of inquiry is never knowledge; it is only ever true belief. See Kvanvig, Jonathan. 2003.

¹⁷⁴ Cases of insignificant knowledge and true beliefs already challenge that assumption.

believed it. If so, then it would be better for Jones to merely believe that p as opposed to knowing it. Confidence may not always accompany knowledge, but sometimes it does. Would it be better for Jones both to know that p and not be overconfident—which causes him to make sloppy inferences? Of course. But Jones is a flawed agent. And sometimes he will get better results from the role a true belief plays, which in this cases does not cause him to be careless.

Or take a case in which an agent is best off having no belief about p , but given she cannot help but form a belief that p , she is still better off not *knowing that p*. Suppose Cassandra is a medical doctor whose family has a history of Alzheimer's. A doctor has informed her and her siblings that it is possible to test everyone in their family to see if they have the gene that will one day cause them to develop this disease. Her sister felt the need to know if she has the gene. Her brother is undecided if he wants to know. Yet Cassandra is emphatically against knowing. She even advised her siblings not to find out the results. Being a physician, Cassandra knows that finding out the results will only entail negative consequences. Currently, there is no way to avoid or cure Alzheimer's disease. Knowing may be harmful, depressing, discouraging, distracting, or cause a sort of pervasive anxiety. Now should a cure become available, Cassandra will be first in line to be tested for the gene. But given her background in medicine, she knows by experience the problems associated with obtaining knowledge about things that no one can do anything about. There is no reason to be distracted by knowledge about which no one can do anything.

Cassandra is a pessimist, and so she cannot help but form the belief that she has the gene. Cassandra is probably better off not forming any belief, but given she cannot help but form some belief about her condition, she might be worse off if she goes looking for evidence of it.

Suppose a person can only come to know some particular topic by being force-fed. The other option is she exercises her intellectual virtues autonomously and achieves at the most only true beliefs about that topic. She might be better off, epistemically speaking, with the second approach.¹⁷⁵

Sometimes obtaining knowledge requires many more resources than mere true belief. In such a case, true belief may be more valuable since knowledge is either not possible or costs too many resources to acquire. This last point might be challenged on the following grounds. *If* knowledge were possible and *if* it did not require one to sacrifice resources that could be better used elsewhere—to accomplish other epistemic aims—knowledge would still be better than true belief.

I wish to give a few replies to such an objection. First, knowledge and true belief often do not cost the same to acquire.¹⁷⁶ Knowledge often comes at the expense of one's cognitive resources. What is easy for someone to believe may be significantly more difficult to know. Sometimes knowledge is more expensive to obtain, which often varies from agent to agent and environment to

¹⁷⁵ See below the demon and child-raising cases.

¹⁷⁶ Riggs acknowledges this point.

environment. One may have great difficulty obtaining knowledge that p at time t_1 , but then come by it more easily at t_2 .

Second, the objection assumes knowledge and true belief can be meaningfully evaluated discretely. Kvanvig rightly points out that any evaluation of discrete beliefs—and therefore knowledge—is mistaken since beliefs cannot be understood apart from their relation to other beliefs. So we cannot simply evaluate them as if they were isolated or cut off from the rest of an agent's beliefs and other epistemically relevant properties.¹⁷⁷ Consequently, it is questionable how helpful the conclusions are that one develops by isolating individual beliefs. It might be true that, given the cost of obtaining $K(p)$ and $TB(p)$ are equal, knowledge is better than $TB(p)$. But the conditional is often false. So the conclusion is not very interesting; it is conditionally true. And that condition often does not obtain. So it is false that knowledge is always better than true belief.

These types of conditions—that assume getting knowledge is never more intellectually demanding than acquiring mere true belief—seem to me so grossly abstracted away from real world people that they appear unrelated to *human-centred* epistemology. When we analyse discrete beliefs and set the conditions, costs, and payoffs of belief acquisition and knowledge to ineffectual or nil, we

¹⁷⁷ Kvanvig, Jonathan. 1992. 181-182.

are doing epistemology for some other species. And why should we care about *that*?¹⁷⁸

But suppose we grant the implausible conditions that knowledge always has the same payoffs and never uses more resources that could be spent elsewhere to achieve other epistemic ends. Would that save the assumption that knowledge is always more valuable than true belief? If there are truths with no value then it will not. Truths like the n number of motes of dust on my desk appear to have no value whatsoever; they are utterly useless. But if the truth of the number of motes on my desk has no value, then it is not clear how knowing that would have any value either, even if we grant the conditions above. Cases involving nosiness seem to me in a similar boat. Even if it did not take more resources to gain knowledge, some information that we identify as gossip is either not worth knowing or even warrants blame for knowing.

A critic might point out that it is not truth, but rather *true belief* that has value.¹⁷⁹ So we might say this about the value of true belief:

<EG>If I have a belief that p , $B(p)$, I want p to be true.

The strategy here is to claim that there is, even if minimally, some value in having a true belief. And if so then one might think there is more value in knowing rather than simply true believing, regardless of how trivial those true beliefs are.

¹⁷⁸ Despite epistemologists' claim to have acknowledged a more realistic human agent, it seems to me there is always a threat in the literature of forgetting what species we are theorising about.

¹⁷⁹ Sosa, Ernest. 2004.

Unfortunately, <EG> is false. <EG> has a rather obvious problem: desire does not entail value. Many things are desired; not all of them are valuable. This is one of the lessons we have learned from expressivist ethics. Desire and value are not coextensive. <EG> then must be understood in relation to the virtuous agent. That is:

<EG²> If a *virtuous agent* $B(p)$, she wants p to be true.

But <EG²> is still false. Virtuous people do not always want their beliefs to be true. I may for example, believe my wife has had a miscarriage, but want it to be false. So we can still improve upon the desire for true beliefs.

<EG³> A virtuous agent wants to $B(p)$ if and only if $\langle p \rangle$ is true.

This appears to fix the problems above. The goal of the virtuous agent is simply the truth goal identified in chapter one. Understood this way the critic attempts to attribute value to trivially true beliefs. Successfully salvaging the third assumption turns on the veracity of <EG³>. But what reason is there to think <EG³> is true? I think there is good reason to reject <EG³>. ¹⁸⁰

Epistemic exemplars often do not care about whether their utterly useless beliefs are true or false; and I do not see a reason why they should. Suppose in the afterlife Aristotle is talking with God and asks him the following question: "OK, what beliefs did I get wrong? I have eternity here so don't leave anything out." God organises the list of false claims by category: metaphysics,

¹⁸⁰ Sosa has pointed out, rightfully I think, that even if <EG³> is true, that does not mean that truth *as such* is valuable. Sosa, Ernest. 2001a. 49.

biology, epistemology, causation, and so on. After discussing at length several beliefs, God gets to the category we might call utterly useless true beliefs: number of motes of dust on his sandals, grains of sand on the beach, his false memory of the weather on a particular day, and so on. I think it is quite reasonable for Aristotle, who has an eternity, to stop God at this point and object, "Well I don't care about *those* beliefs". But why not? Well, because those beliefs simply are not important in any way, whatsoever. And I do not think we would find fault with Aristotle for not caring about correcting those false beliefs, even though he has an eternity to do so and they may come to him at no cognitive cost.¹⁸¹

Consider counterfactual truths. $\langle EG^3 \rangle$ does not exclude them. They are often quite important. But why should virtuous agents care about all counterfactuals? Surely Aristotle is not intellectually blameworthy if he interrupts God when he starts to list all of the counterfactuals Aristotle got wrong or had not even considered.

Or suppose Einstein was bored one day while waiting for a colleague and started to count the number of hairs in his moustache. He forms some belief about the total number of hairs he has just counted. I do not think Einstein would have cared later on if someone came up to him and not only let him know his belief was wrong, but also offered to tell him the truth about the number of hairs in his moustache. Nor is it clear he should care. It does not seem to count

¹⁸¹ He is getting it for free in a sense; God is telling him for nothing.

against Einstein's virtuous intellectual nature if he simply shrugged his shoulders upon hearing the news of his false belief as well as the offer to correct it. In short, it is false that virtuous agents always have a preference to believe things that are true and only true. Sometimes they simply do not care.

<EG³> is dubious. And so it will not do as a strategy to identify value even in trivial cases of true belief. Even if it could, that alone would not be sufficient to show that *knowing* minimally valuable truths is any better than simply believing them. One will also need to buy into the cost-payoff condition above—that knowledge and true belief have equal costs and payoffs. The third assumption is difficult to salvage.

One upshot of this discussion is that the question in Plato's *Meno* is misleading. If I am correct, there is no single value problem but rather many:

1. Why is knowledge sometimes more valuable than its subparts?
2. Why is knowledge sometimes no more valuable than its subparts?
3. Why is knowledge sometimes less valuable than its subparts?

Interestingly, once one has acknowledged value pluralism—as Riggs, Sosa, and Zagzebski claim to be sympathetic toward—the above added value problems are to be expected. The reason is quite simple. If truth is not the only epistemic value; that is, if some values cannot be reduced to truth, then surely there are going to be cases in which values conflict; sometimes a person is only able to secure one value at the cost of another. The very presence of one value might be a hindrance for obtaining another value. Sometimes succeeding in one

aspect of the truth goal requires failing the other. What virtuous agents do well is balance these goals. They recognise that obtaining true beliefs sometimes requires exposing one to acquiring false ones. In fact, as we have seen, sometimes agents arrive at a particular true belief only because they made an inference from false ones.¹⁸² So even with this sort of monism¹⁸³ we can see that value sometimes precludes value.

Agents must discern which values are worth going after in particular contexts. Exemplars will often go after true beliefs. And sometimes the truths they seek are of such a nature that it warrants knowing them. Other times virtuous agents will take up an attitude of epistemic indifference. Even in cases of valuable truth, there might be other epistemic values in play that threaten or override the value of knowledge. One must discern when knowledge has less value than merely true belief. Value pluralism anticipates and explains the existence of the value problems.¹⁸⁴

¹⁸² For example, we simplify scientific explanations to include literally false statements for the purpose of generating true beliefs. Teachers sometimes use this strategy when explaining particularly difficult ideas or theories to students who are not sufficiently able to grasp more complex truths.

¹⁸³ Or dualism if one prefers.

¹⁸⁴ Of course, this is not logically entailed by value pluralism. We can imagine a possible world in which the values never conflict. But that is not our world. Likewise, there is a possible world in which the two aspects of the truth goal never conflict either; but that also is quite far removed from the actual world.

Understanding

Virtue epistemologists have become increasingly aware of the absence of discussions on understanding. Theory of knowledge, particularly as it relates to scepticism and more recently the Gettier problem has been much of the focus. One might have expected a different history considering that both Plato and Aristotle were interested in understanding, arguably more so than in propositional knowledge.¹⁸⁵ In fact, some scholars have argued that the word *episteme* in Plato's *Meno* is more appropriately translated "understanding", not knowledge.¹⁸⁶ Leaving exegesis aside, virtue theorists are warranted in exposing the parochial nature of the history of epistemology. Given the primary object of discussion has been propositional knowledge and attempts to distinguish it from mere true belief, it should not surprise us that epistemologists sound like value monists. The volumes of work done to develop and refine concepts of justification, warrant, reliability, proper function, and so on served mostly to create an environment that restrained and reinforced this myopic interest. Those who did consider issues outside propositional beliefs were often thought to be doing "broad epistemology", as opposed to epistemology proper.¹⁸⁷

Once the enterprise of epistemology is described in such a way, it is natural that philosophers not only look like value monists but also eventually

¹⁸⁵ Zagzebski points out that both Plato and Aristotle discuss the relationship between understanding and *techne*, which is sometimes translated as craftsmanship or art. Zagzebski, Linda. 2001b. 238-241.

¹⁸⁶ See Zagzebski, Linda. 2001b.

¹⁸⁷ DePaul, Michael. 2001.

take it for granted. Taking propositional knowledge to be the central concern, one is in the business of describing the status of discrete beliefs, specifically, their truth value. However, a value-driven approach construes epistemology more generally as the discipline involving belief-acquisition, maintenance, revision, inquiry, and intellectual value, while leaving open what roles concepts such as truth, knowledge, and understanding play.¹⁸⁸ Fortunately, philosophers have become increasingly interested in understanding and its place in epistemology.

I want to make the three following claims about understanding.

1. Understanding belongs as an area of study to epistemology.
2. Understanding is often epistemically valuable.
3. Virtuous agents often aim for understanding.

I suspect few epistemologists would challenge those claims. I will not develop a theory of understanding here. Nonetheless, I want to make a few observations that imply value monism is unlikely to account for the role of understanding.

If it turns out that the value of understanding cannot be captured entirely by appealing to propositions and their truth values, value monism is in trouble. I will argue that although propositional content and truth can contribute to one's understanding, understanding is more than their combination. In fact, I suspect

¹⁸⁸ The Stoics, for example, appeared more concerned with mastering epistemic skills and practical wisdom than propositional knowledge. Annas, Julia. 2003. See also Riggs, Wayne 2008. And Hookway, Christopher. 2003.

part of an agent's understanding often includes false beliefs. First let us distinguish between understanding and knowledge, then more generally propositional content.

This distinction dates back to at least Plato. According to Zagzebski, Plato thought understanding was gained by one's skill or mastery of *how to do* something well and might have even required non-cognitive processes, perhaps actions. Plato identified a close relationship between *techne* and understanding, where *techne* stands in contrast to *episteme*, which is typically translated *knowledge*.¹⁸⁹

Not only is there some historical precedence for making the distinction, but there is also a commonsense reason reflected in ordinary ascriptions. Take the way in which we might distinguish between two agents with regard to some area of study. We differentiate between one who has collected true beliefs—and knowledge—about a topic or body of beliefs and one who exhibits a deep understanding of it. For example, we can distinguish between someone who is knowledgeable about a culture—perhaps she knows many facts and trivia—and one who also understands the culture. She *understands* the facts and trivia in a way the other person does not. She has some insight, can manoeuvre, apply, and appreciate the knowledge she has of that culture. This understanding may have come in part by sharing many beliefs of the culture, some of which are false.

¹⁸⁹ Zagzebski, Linda.2001b. 241.

This ordinary distinction comes from, I think, our recognition that there is a difference between knowing a belief or set of beliefs to be true and having a particular insight into the *relationship* between those true beliefs and how they *relate* to the world. When someone is recognised to have, not only knowledge, but also understanding she is thought to have achieved a sort of pattern recognition. Kvanvig calls this an “appreciation” of how various elements of a body of information are related to each other in terms of explanation, logic, probability, as well as other kinds of relations.¹⁹⁰

Notice that the object of understanding is a *body* of information and the *relationship* between beliefs—not the beliefs themselves.¹⁹¹ To see this, consider a single proposition. Two people can equally know a proposition to be true yet have different degrees of understanding. Likewise, I may know some proposition for a long time and only after years of knowing it come to have a deeper understanding of it.¹⁹² Or suppose I know there was a car accident outside my house because my wife told me she saw it happen only minutes ago. My curiosity may cause me to go outside to see the accident for myself. However that will not increase my knowledge that “there was a car accident”. I already knew that. My curiosity is directed at gaining some understanding that I lack. It is true of course that I will pick up more knowledge as I seek to satisfy my curiosity. But I do this only to deepen my understanding of what I already know.

¹⁹⁰ Kvanvig, Jonathan. 2003. 192-193. Roberts and Wood also call this epistemic appreciation. Roberts, Robert and Jay Wood. 2007. 33.

¹⁹¹ Zagzebski, Linda. 2001b. 217.

¹⁹² Roberts, Robert and Jay Wood. 2007. 33.

One's degree of understanding corresponds to one's sensitivity and appreciation of the relation between beliefs. There is nothing awkward about ascribing degrees of understanding. We naturally think of understanding in gradation. With notable exceptions, many philosophers find it somewhat awkward to think of propositional knowledge in degrees.¹⁹³ People either know some proposition, p , or they do not. A few exceptions to this include contextualist and contrastive accounts of knowledge as well as Hetherington's views. Regardless of one's theory, some items of knowledge appear incompatible with degree.¹⁹⁴ Yet it is quite natural to wonder how much one understands some proposition p . And I often seek to deepen my understanding of p , where p is judged to be valuable or worth understanding. We might debate whether entertaining defeaters for p can increase one's knowledge that p . But no such discussion is required for understanding.

There are other differences to note between understanding and knowledge. First, as both Kvanvig and Riggs¹⁹⁵ have pointed out, understanding and knowledge differ in the way they are undermined. For example, one's understanding of a topic, object, or person may involve several false beliefs. In fact, we often help people to understand a topic by intentionally promoting false beliefs;¹⁹⁶ my understanding of quantum mechanics consists of several false beliefs. This is not untypical of scientific explanations, which include statements

¹⁹³ Kvanvig, Jonathan. 2003. 195 -196.

¹⁹⁴ For example, simple math and logic.

¹⁹⁵ Kvanvig, Jonathan. 2009.

¹⁹⁶ Zagzebski, Linda. 2001b. 244

that are often false.¹⁹⁷ But that does not undermine one's understanding of the topic, object, person, or phenomenon in question. This is also true of models, which are not literally true. We often use models for explanatory purposes. That the model itself is not true does not undermine the understanding one gains from it. This is because models are not the objects of understanding, just as true and false statements in scientific explanations are not. Instead, how a model *relates* to the world is the object of understanding. By contrast, one's knowledge *is* undermined by false beliefs. Gettier famously showed that.¹⁹⁸

Gettier-type considerations also teach us that while knowledge is undermined by salient luck, understanding is not. For example, I do not know that I am pointing at a real barn if it is simply a matter of luck that I did not pick out one of the many fake ones in the area, which are all indistinguishable to me from real barns. Likewise, I cannot be said to know that it is three o'clock if the clock I am looking at is—unknown to me—broken and just happens to indicate the correct time. Luck plays no such role in one's understanding. Suppose I am going on vacation to China where there are substantially different cultural norms than what I am familiar with in Canada. So I seek a minimal understanding of the customs by reading a guidebook on Chinese customs. But suppose the guidebook was written by someone who is not only lying about being an expert on Chinese culture, but compiled her guidebook by visiting random websites without

¹⁹⁷ Kvanvig, Jonathan. 2009. 342-343.

¹⁹⁸ Gettier, Edmund. 1963.

checking any of their credentials. Some of the sites are notoriously unreliable and riddled with false information. And some are nothing more than compilations of stereotypes. So the guidebook was written carelessly to say the least. But as luck would have it, everything written in the book is true. Having read the book, believing it to have been written by an expert, I go to China and manage to follow their cultural norms flawlessly while I am there. I am even able to explain them to other Canadians should they ask. Despite the significant amount of salient luck, surely I still understand Chinese culture.

Finally, potential defeaters undermine knowledge but not understanding. To see this we can simply adjust the above case of the lying expert. Suppose the author has previously lied about being an expert on other cultures so she could write a variety of guidebooks, which turned out to contain little more than stereotypes and inaccurate accounts of cultural norms. Let us assume her previous books were carelessly written. They received terrible reviews by real experts, who exposed this author as a fraud. However, because of the terrible reviews, her latest book on Chinese culture was carefully researched and she ended up producing a rather accurate guide. I pick up the book and read it in preparation for my trip to China. Unknown to me, however, the expert critics recognise the author of the guidebook and have unanimously given negative reviews. None of them recommend the guidebook; they warn vacationers to stay far away from this author's work. We can even suppose this negative assessment of the guidebook is reported in a segment on the news covering guidebooks. I

am aware of neither the testimony of the expert critics nor the news report. Nevertheless, I go to China and manage to follow their cultural norms flawlessly while I am there. Again, I am even able to explain them to other Canadians should they ask about Chinese culture.

I do not think there is any question about whether I have some understanding of Chinese culture. However, given the presence of potential defeaters—the expert reviews and news reports advising people to stay away from this author’s work—my beliefs about Chinese culture are not knowledge. So the conditions for undermining knowledge and understanding are different.

I conclude here by claiming that knowledge is not required for understanding. Nor is understanding required for knowledge. Since understanding and knowledge are independent, understanding is often valuable, and there is no obvious pattern in terms of knowledge when understanding is more or less valuable, understanding has its own value.

Autonomy and Responsibility

I want to conclude this chapter by quickly discussing a few aspects of virtue epistemology I think are better handled by pluralism. First, consider epistemic autonomy and responsibility, both of which are highly valued. Many virtue epistemologists point out that what is so valuable about knowledge is at least in part that the true belief in question is being obtained *through the agent’s*

performance such that it is to her credit. When understood this way, it presupposes what we might call doxastic autonomy. A person is doxastically autonomous when she has an appropriate amount of freedom in gathering information, entering deliberation, reflecting, and evaluating the quality of that information to form, maintain, and revise beliefs.¹⁹⁹ The opposite is something along the lines of brainwashing. The amount of freedom required for doxastic autonomy is not something we need to worry about at this moment. I only want to claim here that there is an intelligible distinction between someone who has formed her beliefs with a degree of autonomy at some salient point in the process and one who has not—perhaps due to features of the agent’s environment, which may include some form of excessive influence, coercion, and lack of cognitive capacities. We can distinguish doxastic autonomy from epistemic autonomy in that the latter involves more than freedom governing one’s beliefs. It includes freedom to develop other mental states such as understanding, appreciation, and acquaintance, which increase one’s epistemic status or condition. For example, someone who understands *x* is often better off, epistemically, than someone who does not understand *x*.²⁰⁰ Consequently, a person is epistemically autonomous when she has a sufficient amount of

¹⁹⁹ This does not presuppose doxastic voluntarism.

²⁰⁰ Assuming *x* is not epistemically harmful or trivial.

freedom to shape her epistemic status, which consists of more than sets of beliefs. Related to autonomy is epistemic responsibility.²⁰¹

Both autonomy and responsibility are thought to be valuable. However, it is dubious they can be accounted for by the monist. According to monism, the only epistemic value is the veristic one. But a brainwashed agent not responsible for her beliefs can have just as many true beliefs as an autonomous one. Or consider being offered a chance to acquire a significantly high ratio of true beliefs but at the cost of epistemic freedom and responsibility. A shadowy figure is offering to sit an agent down in a chair and force images, thoughts, and beliefs into that agent such that it will produce a higher ratio of true beliefs than if she were to retain her autonomy and responsibility. Would a virtuous agent always accept such an offer? I highly doubt it. I know I would not. That is not to say there are *no conditions* under which I would accept the offer. If, for example, the truths were particularly important ones about the world that I could not otherwise acquire, I may accept the offer. Likewise, if the topic was useful but boring to learn, I may opt to have the information forced upon me. But knowing only that the end result is a higher ratio of true beliefs is not a sufficient epistemic payoff to conclude that a virtuous agent will always accept such an offer. Epistemic autonomy and responsibility are too valuable to sacrifice for a high ratio of true beliefs. And I do not see that I would be committing any

²⁰¹ The relationship is quite controversial. For example, not everyone agrees autonomy is required for responsibility.

epistemic shortcoming here if I were to avoid brainwashing. Notice that for the monist the decision should be easy. If truth is all one cares about, there are no other worthwhile considerations. One's decision might be influenced by domains outside of epistemology. But insofar as we are only concerned with the epistemic domain, monists are forced to the awkward position that it would be more intellectually virtuous to accept the brainwashing since, on their view, autonomy and responsibility are valuable only because they are thought to lead to a higher ratio of true beliefs. That is not true in the above case, so there is no reason to reject the offer of brainwashing. As a monist, one might even have a *duty* to accept it. This is a sort of swamping problem. The value of autonomy and responsibility is being swamped by the value of true belief. A pluralist is not forced to such a conclusion and has no swamping problem because she can recognise that several values are in play; epistemic autonomy and responsibility are perhaps irreducibly valuable.

A similar example is raised by Casey Swank. A demon offers us a chance for lots of true beliefs in a demon world, but at the cost of being utterly close-minded, unreasonable, and dogmatic.²⁰² Would we take the offer? I suspect not. But if not, then we identify some virtues and vices independent of their relation to true beliefs.

There is nothing bizarre or extraordinary about the demon case. Every parent must make a similar decision when raising a child. Suppose you have

²⁰² Swank, Casey. 2001. 200-202.

children. You can indoctrinate them and if you do it effectively they will adopt many of what you consider to be true beliefs. Or you can make open-minded inquirers of them, in which case they are likely to disagree erroneously with you on many views. What do you want for them?²⁰³

Autonomy and responsibility are valuable. Sometimes we think less of someone—intellectually speaking—who has a true belief because she got it cheaply or in a lazy way, perhaps she looked up the answer or referred to an expert instead of getting the answer by using her own cognitive abilities. This is not always the case of course; sometimes we praise someone who did not bother using resources unnecessarily. Short cuts can be both praised and blamed; the different evaluations depend on the context. In one context we interpret an agent's actions to manifest a sort of intellectual laziness. In other environments, short cuts are an efficient use of one's resources. Interestingly, none of these considerations should matter to the monist.

Why not? Let me make the point by two analogies. Suppose that all you care about is money.²⁰⁴ If so, then it does not matter how you get money as long as you get it. It does not matter if you won the lottery, found it, earned it, stole it, worked harder, were smarter, and so on. What matters is the money. If however, someone comes along and claims only to value money, but then thinks that earned money is better than stolen money, we should conclude that this

²⁰³ Adam Morton pointed out to me the analogous case of child-raising.

²⁰⁴ DePaul, Michael. 2001. 175-179.

person values more than money. Money might remain this person's *highest* value, but it is not the only one.

Likewise, someone who prefers not only to win a game, but also earn it against a capable opponent values more than winning. For example, placing first in a tournament might be an athlete's highest value, but the victory will be sweeter if she wins without her opponents forfeiting to her or playing injured. For many, *enjoying* the competition is also important. Such people are not athletic value monists. They value more than the victory.

The lesson, I claim, carries over to true beliefs. If true belief is the only valuable thing, it does not matter how one gets the true belief, or what is sacrificed to get it, epistemically speaking. But virtuous agents prefer autonomy and responsibility to accompany their belief acquisitions, whether true or false. In fact, the child-raising example implies that sometimes an autonomously acquired false belief is more valuable than a force-fed true one. So monism is ill-equipped to account for epistemic autonomy and responsibility.

Stupidity

Virtue epistemologists tend to focus their attention on virtues. However, an examination of epistemic vices opens up the possibility of identifying new values. Take for example the evaluation of someone who is described as being stupid. What is meant when such a thing is uttered? I suggest this ascription

amounts to more than saying that the person in question has a small ratio of true beliefs. That might fit the description of what we mean when we say someone is ignorant, especially ignorant of a topic.²⁰⁵ But when we evaluate someone as being stupid—whether the evaluation is warranted or not—we are saying more about that person’s ability to understand, make skillful judgements, employ practical reasoning, and generally navigate the world around her. And it seems possible to get, stupidly, true beliefs.

This view is supported by our recognition that people sometimes place themselves in environments that are beyond their competency. So we might say of someone in a particular context that she is epistemically incompetent. Surely this is not a description of any ratio of beliefs—even relevant ones—but rather her capacity to understand and appropriately make her way through her current environment. I take stupidity to be a generalisation of one’s incompetence. A person is stupid when she has a wide range of environments in which she is incompetent. To be clear, a token ascription of stupidity may involve a lack of ability to acquire true beliefs, but this does not exhaust the vice altogether. Sometimes we identify a person as being stupid because she fails to appropriately *apply* or *appreciate* the true beliefs she already holds.²⁰⁶ It sometimes includes an evaluation of the *relation between* beliefs, the world,

²⁰⁵ Notice that the evaluation of ignorance is not simply a high and low amount of false and true beliefs respectively. Rather, it is of *particularly salient* beliefs.

²⁰⁶ Our evaluations recognise a relation between competency and autonomy. As Descartes warned, if one is well beyond their competency limit, exercising autonomy may be undesirable. In fact, it is a sign of intelligence when a person knows that she is out of her depth and purposely limits personal freedom.

further inquiries, and decision-making. Monism cannot give an adequate account of stupidity since it is not exhausted by appealing to the truth goal. Negative evaluations such as stupidity imply pluralism.

Conceptual Priority and Community

The following remarks will serve mainly to tie together some points previously made. Virtue epistemology is supposed to conceptually prioritise the agent. Unfortunately, monism threatens to make this priority somewhat superficial. On a monistic account, virtues are defined by their relation to true believing. But then true believing must be conceptually prior to virtuous character. One cannot even identify intellectual character without first appealing to true believing.²⁰⁷ Or to use Simon Blackburn's language, you would not even know what game is being played unless you were prioritising truth, having determined the sole intellectual value before the game begins. If one defines intellectual virtues only by their relation to a prior evaluation of beliefs, one is no longer a virtue epistemologist; the entire discussion of character traits presupposes a prior—and solitary—value. Blackburn's own solution for monists is to adopt a deflationist account of truth. If a deflationary view of truth makes the pluralist uneasy, she need not worry since on her account virtuous epistemic living is not exhausted by the truth goal. So she has no such problem.

²⁰⁷ Blackburn, Simon. 2001. 22.

Jonathan Dancy has pointed out that some virtues are identified *before* their relation to the truth goal is considered. A candidate suggested by Dancy is intellectual tolerance. It is not clear that tolerance is conducive to the truth goal; but we do not hesitate to ascribe this virtue to people who display it. Now, we might be able to show after the fact that such a virtue is truth-conducive, but it is not built right into the virtue itself.²⁰⁸ Some virtues might be defined in terms of the truth goal, but not all of them. When one makes a shift in conceptual priority, from the truth value of beliefs to the character of agents, this should be expected.

Prioritising the agent means acknowledging her cognitive powers and relation to the environment. I want quickly to revisit one aspect of the environment because it plays such a salient role in shaping the agent's telos: community. Even an isolated meditation such as one done by Descartes is embedded in a particular tradition.²⁰⁹ Scientists work within socially developed and guided research projects. And as anyone who has worked in a group knows, when we work to solve problems, we do not simply add collective cognitive powers. We consider the group make-up. Does it include people who are better at the fine details? Best to use them accordingly. To be most efficient, a group is going to require that some of its constituent parts—individual agents—perform in ways that might not be efficient for individual goals and interests. It might

²⁰⁸ Dancy, Jonathan. 2001. 79.

²⁰⁹ Roberts and Wood. 2007. 114.

require a person to solve problems in areas of her weakness. Some people's skills may be at odds and so individuals may need to compensate. Heuristics need to be adjusted depending on the social context and composition. This may even entail individuals giving up truths so that others can obtain them. Consider, for example, a father's attempt to help his son understand a math assignment. While the father may be in over his head and acquire many false beliefs, he may grasp enough just to get by and help his son learn.²¹⁰

Consider a sports analogy. The best performance of an individual hockey player is going to depend greatly on the capacities of her teammates. Suppose player *x* has the puck and is close to the opposing team's net. Should she take a shot at the net? Well, that depends on both the positioning of her teammates as well as their relative skills at puck handling and shooting. If player *x* is on a breakaway with no teammates in sight, she is for all practical purposes, no longer playing on a team; the game can be reduced in that moment to player *x* against player *y*—the opposing team's goaltender. In those cases, shooting is the obvious answer; that is the only way to get the goal. However, once we acknowledge the presence of teammates and their skills, it is sometimes more appropriate to refrain from shooting and instead pass the puck. Good hockey players judge when to pass and when to shoot. It makes little sense for a hockey player to shoot on the net if her teammate—who is generally considered to be an excellent goal scorer—is in an advantageous spot to get a goal. Likewise, it

²¹⁰ Adam Morton put forth this example.

would be poor judgement to pass the puck to a less skilled player with no line of sight toward the net, especially if the player with the puck has an opportunity to score.

I suspect we are rarely on a genuine epistemic breakaway. Improving our epistemic status is usually embedded in a social context that has its own interests and goals. Just as it makes little sense to advocate one ideal strategy in hockey—shooting on the net—it makes little sense to think virtuous agents are always aimed at true beliefs. As we saw above, the goal of the game can be reduced to one thing: winning. But the manner in which individuals win, especially once we acknowledge they are not playing a solo sport, varies by context. The winning strategy does not consist of every player trying to score. The father helping his son with the math problem is a team player. He passes in order to give his son the chance to score. The telos of a virtuous agent depends in large part on her community, an insight achieved by taking seriously the view that agents and their virtues ought to be considered apart from only one intellectual value.

Conclusion

One of my aims is to widen the scope of virtue epistemology. A value-driven approach is well positioned to do such a thing. Instead of taking for granted that the truth goal is the only epistemic aim, we ought to be open to the possibility that there are other epistemic values, some of which are incompatible with acquiring token truths and avoiding false ones. As I have argued, once we look at our epistemic evaluations, it is clear we are concerned with more than propositional knowledge, true belief, and error avoidance. Our intellectual lives are far too complex to have such a myopic view of agent-based epistemology.

It may strike the reader as suspicious that despite urging a wide scope of epistemology, I have apparently neglected one of the most important aspects on the topic: wisdom. Wisdom, is not only epistemically praiseworthy, but arguably the highest possible epistemic achievement. Whatever wisdom amounts to, it would be difficult to deny that wisdom is—or ought to be—a central issue in epistemology. Consequently, I propose that any theory of virtues that does not have a central role for, or appropriate account of, wisdom is problematic. But then why have I neglected it so far? To the reader I reply: wisdom has been pervasive in the discussion and neglected in name only.

I have argued that epistemic virtues are unified by their relation to the telos of the virtuous agent, that is, the epistemically well-lived life. I have tried to describe this as intellectual flourishing, good sense, epistemic success, and a kind

of fitness where the agent is able to navigate her environment successfully, solve problems, and inquire into the nature of world. While none of those expressions is exhaustive, they move us away from mistakenly thinking the telos can be described by appealing to the ratio of discrete or time-sliced beliefs. The above expressions give a partial illustration of the goal of virtuous agents. Arguably, this telos can be summed up in one word: wisdom. The term 'wisdom' comes with much historical baggage. So while I do not insist here that it is the appropriate term to identify with the goal of the virtuous agent, I think it is a good candidate. Wisdom does not look like any one thing. It is not a solitary ideal that can be narrowly evaluated in every agent. What it means to be wise will be somewhat different from agent to agent. With that in mind, my view can be summarised in the following way: The goal of virtuous agents qua intellectual beings is to be wise. It is plausibly the highest epistemic aim.

The following observation, properly understood, is a rephrasing of what I have already written. Wisdom cannot be accounted for by value monism. Whatever the details of wisdom turn out to be, it likely includes some experiential and practical aspects. It involves having good judgement, for example. Perhaps a god can achieve such a state from an armchair. However, insofar as we are interested in human-centred epistemology, wisdom requires some fieldwork. Good judgement comes from experience, which as it turns out, often comes from bad judgement. It would be quite difficult to remove the experiential and practical aspects altogether from wisdom. Sometimes, though

perhaps not always, we identify a person's wisdom *because* of her skill in appreciating, applying, inferring, and using her beliefs, often through some action. Being wise is not simply a matter of having more true beliefs.

Similar remarks can be made about cleverness. An agent may be described as clever when she exhibits, for example, a unique approach to a topic or problem.²¹¹ Cleverness is not identified by evaluating the truth value of the agent's total set of beliefs. Yet it is an epistemic achievement. A view that denies a plurality of values will have a difficult time trying to account for epistemically prized states such as wisdom and cleverness, both of which involve an agent *doing something* with her beliefs.

Some virtue epistemologists worry that what an account like mine is doing is conflating practical and epistemic domains. Sure there are epistemic issues that are also practical, says the critic, but the two are not one and the same. So we should retain a distinction between practical and epistemic value. The critic might worry that a sort of pragmatic encroachment is threatening the epistemic.

I want to make it clear that on my view the pragmatic value of a belief is distinguishable from its epistemic value. However, it would be a mistake to think that practical and epistemic domains are always mutually exclusive. We cannot always drive a wedge between them. Sometimes the practicality of one's inquiry

²¹¹ Likewise, we describe an agent as being "sharp" when she thinks quickly. This seems to me both epistemically valuable and not reducible to the truth goal.

is necessary for a positive epistemic evaluation. Once virtues are defined in part by indexing them to environments, practical concerns are unavoidable. This follows from the very nature of virtues and one of the motivations for their introduction in ethics. If one is interested in describing epistemic flourishing, one better start talking about *life*.

Imagine in heaven, prior to being born on earth, a salesman is tasked by God to allocate epistemic virtues. Jones gets to pick out a few. Jones has not seen earth, nor does he know his physical talents, learning style, social relations, and resources that will be available once he is born. He does not even know what it is like to be in a body. Suppose the salesman asks Jones to pick some virtues and he declines the offer. How will the salesman pitch his product? What can he say about the virtues that will pique Jones' interest if he cannot appeal to *anything* practical? The salesman will need to appeal to some possible experiences in life, for example, some environments, problems to be solved, methods of gathering information, relations and freedoms within various societies. Without presupposing these backgrounds, the entire notion of "virtue" becomes unmotivated, deflated, and perhaps even meaningless. I suspect Jones would not even comprehend the product being offered to him. This is especially noticeable when we consider virtues of limitation management. All of them are essentially practical. Remove the practicality and you will have removed the virtue.

Consider what it means to get the lie of the land. This consists of knowing how to handle a certain environment, including social ones. We might get a rough idea of a topic so that we can make our way through it for our purposes. Other times we get the lie of the land specifically so that we do not need to think in order to navigate it. These are essentially about staying in contact with one's environment. They are practical and not always aimed at the truth goal. Nonetheless, they are epistemic.

We should be suspicious of any attempt to remove life from epistemology, which is what one is doing when one tries to carve off the intellectual from the agent's relationship to her environment.²¹² Some intellectual virtues are *defined* by their practicality, never mind that virtues more generally are defined by context. That does not mean we need to define epistemic value entirely in terms of usefulness—because practical value is not the only one. Nonetheless, something can be both epistemic and practical; and sometimes they are identified as epistemically valuable *because* they exhibit some practical skill. In such cases, you cannot remove the life, usefulness, and practicality without annihilating the intellectual value. Yet this is what one must do as a monist since, according to that view, values other than the diachronic one are essentially non-epistemic. Epistemic value pluralism has no such problem.

²¹² This seems to presuppose a sort of dualism.

Despite the seeds of pluralism found in Sosa, Zagzebski, and others, I think it is time virtue theorists start taking pluralism seriously and widen the scope of epistemology. They can start by developing a virtue theory that is not cashed out in terms of the truth goal alone. Some candidates of epistemic value include: true beliefs, avoidance of false ones, empirically adequate beliefs, knowledge, understanding, epistemic autonomy, epistemic responsibility, and wisdom. They are valuable because they contribute to the epistemically well-lived life. Only by acknowledging the plurality of values can we begin to form a picture of our multifaceted intellectual lives and how to make them virtuous.²¹³

²¹³ I want to thank Adam Morton for the many conversations, suggestions, and criticisms.

Bibliography

- Alston, William. "Concepts of Epistemic Justification." *The Monist* 68 (1985): 57-89.
- . *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press. 2005.
- Annas, Julia. "The Structure of Virtue." *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford University Press: Oxford 2003. 15-33.
- Aristotle. "Nicomachean Ethics." *The Basic Works of Aristotle*. Ed. Richard McKeon. 935-1112.
- Audi, Robert. "Epistemic Virtue and Justified Belief." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 82-97.
- Axtell, Guy. "Virtue Theory and the Fact/Value Problem." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2000. 177-193.
- Baehr, Jason. "Is there a Value Problem?" *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 42-59.
- Blackburn, Simon. "Reason, Virtue, and Knowledge." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 15-29.
- BonJour, Laurence. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press. 1985.
- BonJour, Laurence. "Sosa on Knowledge, Justification, and "Aptness"." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2000. 87-97.
- Brady, Michael. "Curiosity and the Value of Truth." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 265-283.

- Brogaard, Berit. "The Trivial Argument for Epistemic Value Pluralism, or, How I Learned to Stop Caring about Truth." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 284-305.
- Chisholm, Roderick. *Theory of Knowledge*, 2nd ed. Englewood Cliffs: Prentice-Hall. 1977.
- . "The Place of Epistemic Justification." *Philosophical Topics* 14.1 (1986): 47-54.
- . *Theory of Knowledge*, 3rd ed. Englewood Cliffs: Prentice-Hall. 1989.
- Dancy, Jonathan. "Supervenience, Virtues, and Consequences." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 73-85.
- David, Marian. "Truth as the Epistemic Goal." *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press. 2001. 151-169.
- Davidson, Donald. "Knowing One's Own Mind." *Proceedings and Addresses of the American Philosophical Association* 60.3 (1987): 441-458.
- DePaul, Michael. "Value Monism in Epistemology." *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press. 2001. 170-183.
- DeRose, Keith. "Assertion, Knowledge, and Context." *The Philosophical Review* 111.2 (2002): 167-203.
- Descartes, René. "Meditations on First Philosophy." *The Philosophical Writings of Descartes*. Trans. John Cottingham, Robert Stoothoff, and Dougald Murdoch. Cambridge: Cambridge University Press. 1984.
- Dretske, Fred. "Epistemic Operators." *Journal of Philosophy* 67.2 (1970): 10075-1023.
- Driver, Julia. "Moral and Epistemic Virtue." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2000. 123-134.
- . "The Conflation of Moral and Epistemic Virtue." *Moral and Epistemic Virtues*. Eds. Michael Brady and Duncan Pritchard. Oxford: Blackwell Publishing. 2003. 101-116.

- Elfin, Juli. "Epistemic Presuppositions and Their Consequences." *Moral and Epistemic Virtues*. Eds. Michael Brady and Duncan Pritchard. Oxford: Blackwell Publishing. 2003. 48-66.
- Engel, Pascal. "Pragmatic Encroachment and Epistemic Value." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 183-203.
- Elgin, Catherine. "Is Understanding Factive?" *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 322-330.
- Fairweather, Abrol. "Epistemic Motivation." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 63-81.
- Foley, Richard. "The Foundational Role of Epistemology in a General Theory of Rationality." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 214-230.
- Gettier, Edmund. "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121-123.
- Goldman, Alvin. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73 (1976):771-791.
- . "The Unity of Epistemic Virtues." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 20-48.
- Goldman, Alvin and Erik Olsson. "Reliabilism and the Value of Knowledge." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 19- 41.
- Greco, John. "Two Kinds of Intellectual Virtues." *Philosophy and Phenomenological Research* 60.1 (2000): 179-184.
- . "Virtue, Skepticism, and Context." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2000. 55-71.
- ."How to Reid Moore." *The Philosophical Quarterly* 52.209 (2002): 544-563.

- . "Knowledge as Credit for True Belief." *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford: Oxford University. Oxford. 2003. 111-134.
- . "The Value Problem." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 313-321.
- Hájek, Alan, "Waging War on Pascal's Wager." *The Philosophical Review* 112.1 (2003): 27-56.
- Hall, Richard J and Charles R. Johnson. "The Epistemic Duty to Seek More Evidence." *American Philosophical Quarterly* 35.2 (1998): 129-139.
- Hookway, Christopher. *Scepticism*. New York: Routledge. 1990.
- . "Regulating Inquiry: Virtue, Doubt, and Sentiment." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2000. 149-160.
- . "Epistemic Akrasia and Epistemic Virtues." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 178-199.
- . "How to be a Virtue Epistemologist." in *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski, Oxford: Oxford University Press. 2003. 181-202.
- Koethe, John. *Skepticism, Knowledge, and Forms of Reasoning*. New York: Cornell University Press. 2005.
- Kornblith, Hilary "Zagzebski's Virtues of the Mind." *Philosophy and Phenomenological Research* 60.1. (2000): 197-201.
- . "Epistemic Obligation and the Possibility of Internalism." *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 231-248.
- Kusch, Martin. "Testimony and the Value of Knowledge." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 60-94.

- Kvanvig, Jonathan. *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Epistemology*. Maryland: Rowman and Littlefield Publishers. 1992.
- . "Zagzebski on Justification." *Philosophy and Phenomenological Research* 60.1. (2000): 191-196.
- . *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press. 2003.
- . "The Value of Understanding." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 95-111.
- . "Precis of *The Value of Knowledge and the Pursuit of Understanding*." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 309-312.
- . "Response to Critics." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 339-351.
- Lemos, Noah. "Commonsense in Ethics and Epistemology." *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press. 2001. 204-218.
- Lynch, Michael. "The Values of Truth and the Truth of Values." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 225-242.
- McGonigal, Andrew. "Traditional Epistemology Reconsidered: A Reply to Eflin." *Moral and Epistemic Virtues*. Eds. Michael Brady and Duncan Pritchard. Oxford: Blackwell Publishing. 2003. 67-74.
- Montmarquet, James "Epistemic Virtue and Doxastic Responsibility" *American Philosophical Quarterly* 29.4 (1992): 331-341.
- . "An "Internalist" Conception of Epistemic Virtue." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 135-147.
- Moore, G.E. *Selected Writings*. Ed. Thomas Baldwin. New York: Routledge. 1993.

- Morton, Adam. *Limitation Management*. Oxford: Oxford University Press. Forthcoming.
- Oesterle, John. *Treatise on the Virtues*. Notre Dame: University of Notre Dame Press. 1984.
- Paul, Richard. "Critical Thinking, Moral Integrity, and Citizenship: Teaching for the Intellectual Virtues." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 163-175.
- Percival, Philip. "The Pursuit of Epistemic Good." *Moral and Epistemic Virtues*. Eds. Michael Brady and Duncan Pritchard. Oxford: Blackwell Publishing. 2003. 29-46.
- Plantinga, Alvin. *Warrant and Proper Function*. New York: Oxford University Press. 1993.
- Plato. "Meno." *Five Dialogues*. Trans. G.M.A. Grube. 58-92.
- Riggs, Wayne. "Beyond Truth and Falsehood: The Real Value of Knowing that p." *Philosophical Studies* 107.1 (2002): 87-108.
- . "Reliability and the Value of Knowledge." *Philosophy and Phenomenological Research* 64.1 (2002): 79-96.
- . "Balancing our Epistemic Goals." *Nous* 37.2 (2003): 342-352.
- . "Understanding 'Virtue' and the Virtue of Understanding." in *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford: Oxford University Press. 2003. 202-226.
- . "Why Epistemologists Are So Down On Their Luck." *Synthese* 158 (2007): 329-344.
- . "The Value Turn in Epistemology." *New Waves in Epistemology*. Eds. Vincent Hendricks and Duncan Pritchard. Palgrave Macmillan. 2008.
- . "Luck, Knowledge, and Control." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 204-221.

- . "Understanding, Knowledge, and the *Meno* Requirement." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 331-338.
- . "Open-mindedness." *Metaphilosophy* 41.1-2 (2010): 172-188.
- Roberts, Robert and Jay Wood. *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Oxford University Press. 2007.
- Russell, Bertrand. *Human Knowledge: Its Scope and Limits*. New York: Allen and Unwin. 1948.
- Sosa, Ernest. *Knowledge in Perspective*. Cambridge: Press Syndicate of the University of Cambridge. 1991.
- . "For the Love of Truth?" *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*. Eds. Abrol Fairweather and Linda Zagzebski. Oxford: Oxford University Press. 2001. 49-62.
- . "Perspectives in Virtue Epistemology: A Response to Dancy and Bonjour." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 99-110.
- . "Reflective Knowledge in the Best of Circles." *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press. 2001. 187-203.
- . "The Place of Truth in Epistemology." in *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford: Oxford University Press. 2003. 155-179.
- . "Reply" *Sosa and His Critics*. Ed. John Greco. Oxford: Blackwell Publishing. 2004. 275-325.
- Sosa, Ernest and Laurence Bonjour. *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*. Oxford: Blackwell Publishing. 2003.
- Swank, Casey. "Epistemic Vice." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 195-204.
- Weiner, Matthew. "Practical Reasoning and the Concept of Knowledge." *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press. 2009. 163-182.

Williamson, Timothy. *Knowledge and Its Limits*. Oxford: Oxford University Press. 2000.

Zagzebski, Linda. "The Inescapability of Gettier Problems" *The Philosophical Quarterly* 44.174. (1994): 65-73.

---. *Virtues of the Minds: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York: Cambridge University Press. 1996.

---. "Precis of Virtues of the Mind." *Philosophy and Phenomenological Research* 60.1. (2000): 169-177.

---. "Responses." *Philosophy and Phenomenological Research* 60.1. (2000): 207-219.

---. "From Reliabilism to Virtue Epistemology." *Knowledge, Belief, and Character*. Ed. Guy Axtell. Oxford: Rowman and Littlefield Publishers. 2001. 113-122.

---. "Recovering Understanding." *Knowledge, Truth, and Duty*. Ed. Matthias Steup. Oxford: Oxford University Press. 2001. 235-251.

---. "The Search for the Source of the Epistemic Good." *Moral and Epistemic Virtues*. Eds. Michael Brady and Duncan Pritchard. Oxford: Blackwell Publishing. 2003. 13-28.

---. "Intellectual Motivation and the Good of Truth." in *Intellectual Virtue: Perspective from Ethics and Epistemology*. Eds. Michael DePaul and Linda Zagzebski. Oxford: Oxford University Press. 2003. 135-154.

---. "Epistemic Value Monism" *Sosa and His Critics*. Ed. John Greco. Oxford: Blackwell Publishing. 2004. 190-198.