7-14-2010

The Universalizability Criterion of Reason and the Final Value of Knowledge: A Response to Pritchard's Analysis of the Value of Knowledge

Lisa Lynette Cagle
University of Missouri-St. Louis, lisa.l.cagle@gmail.com

Follow this and additional works at: http://irl.umsl.edu/thesis

The Universalizability Criterion of Reason and the Final Value of Knowledge:
A Response to Pritchard's Analysis of the Value of Knowledge

Lisa L. Cagle
B.A., Philosophy of Human Nature (Individually Tailored Education Plan),
Greenville College, 2005

A Thesis Submitted to The Graduate School at the University of Missouri – St. Louis in
partial fulfillment of the requirements for the degree
Master of Arts in Philosophy

July 2010

Advisory Committee

Berit Brogaard, Ph.D.
Chairperson

Eric Wiland, Ph.D.

John Brunero, Ph.D.

THE UNIVERSALIZABILITY CRITERION OF REASON AND THE FINAL VALUE OF KNOWLEDGE:

A RESPONSE TO PRITCHARD'S ANALYSIS OF THE VALUE OF KNOWLEDGE

by

Lisa Lynette Cagle
B.A., Philosophy of Human Nature (Individually Tailored Education Plan),
Greenville College, 2005

A THESIS

Submitted to the Graduate School of the

UNIVERSITY OF MISSOURI – ST. LOUIS

In Partial Fulfillment of the Requirements for the Degree

MASTER OF ARTS

in

PHILOSOPHY

July, 2010

Advisory Committee

Berit Brogaard, Ph.D.
Committee Chair

Eric Wiland, Ph.D.

John Brunero, Ph.D.

ACKNOWLEDGEMENTS

I am grateful to my thesis committee for their helpful comments on the arguments that

follow. I am indebted to Eric Wiland for his intriguing ideas that spark my interest and

inspire my arguments; to John Brunero for his extensive and careful feedback that make

me a better writer than I could ever be otherwise; and to Brit Brogaard, my thesis chair,

for her constant support and enthusiasm for this project.

DEDICATION

I dedicate this thesis to my father and mother whose support and sacrifices have made my education possible.

ABSTRACT

Kant's account of reason – both the practical account and the theoretical account – can give help in addressing the long-standing problems in epistemology of understanding what knowledge is and why it is distinctly valuable. In this paper, I address these particular epistemic concerns in three main parts: 1) I give an overview of Kant's account of practical reason and argue for the primacy of practical reason over theoretical reason; 2) I present an analysis of knowledge from Kant's account of (practical and theoretical) reason and suggest some intellectual counterparts to some of the key concepts in Kant's ethical theory that line up with his account of theoretical reason; 3) I address four major challenges regarding the value of knowledge and present one promising approach to these challenges offered by Duncan Pritchard. I argue that Pritchard's account fails to adequately address the four key challenges, then apply my Kantian analysis of knowledge to the problems to show how it is superior in addressing them. The positive account I offer here is unique in its explicit relying on duty and the Kantian universalizability criterion of reason.

INTRODUCTION

Virtue epistemology has recently received much attention for using virtue ethical theories as an analogue for epistemic virtue. In this way, virtue ethics gives us new tools to approach both new and longstanding problems in epistemology.

This recent work done in virtue epistemology makes obvious that ethics and epistemology have some things in common. There are at least two things that link ethical theories with normative epistemology: 1) both are concerned with normativity (both are subject to the demands of rationality; epistemology is at least concerned with what is 'reasonable' – or justified -- to believe), and 2) agents can be commended or rebuffed both for the rightness/wrongness and goodness/badness of their actions (or willingness to act) given a set of alternatives for action and for the rightness/wrongness and goodness/badness of holding certain beliefs in the face of possible justification or lack thereof. In ethics, the tradition of practical reason asks the distinctively normative question of what one ought to do, or what would be best to do, given a set of alternatives for action. Agents aim to assess and weigh their reasons for action, which are the considerations that stand in favor of or opposed to the alternative courses of action that are available to them. The kind of reason associated with epistemology – what we might term 'theoretical reason' – can also be interpreted as asking a primarily normative question. That is, theoretical reflection asks what one ought to believe. In theoretical reasoning, agents attempt to assess and weigh reasons for belief, which are the considerations that stand in favor of or opposed to particular judgments one might conclude about the way the world is. Both practical and theoretical reasoning, then, are

concerned with normative regulation: practical reason with the regulation of action, and theoretical reason with the regulation of belief.

Viewing the overlap between ethics and epistemology in terms of the normative aims of practical reasoning and theoretical reasoning allows us to ask whether ethical theories other than virtue ethics might lend us aid in solving epistemic problems, too. Perhaps the most significant moral theory with enduring appeal and in the tradition of practical reason is Kant's deontology.

In this paper, I use Kant's ethical theory to give insight into some major epistemic questions. One long-standing project in epistemology is the attempt to understand what knowledge is and why it is distinctly valuable. I attempt to advance this project in three main parts: 1) I give an overview of Kant's account of practical reason and argue for the primacy of practical reason over theoretical reason; 2) I present an analysis of knowledge from Kant's account of (practical and theoretical) reason and suggest some intellectual counterparts to some of the key concepts in Kant's ethical theory that line up with his account of theoretical reason; 3) I address four major challenges regarding the value of knowledge and present one promising approach to these challenges offered by Duncan Pritchard. I argue that Pritchard's account fails to adequately address the four key challenges, then apply my Kantian analysis of knowledge to the problems to show how it is superior in addressing them. I conclude with big picture implications of my analysis, admit some of the limitations of my work, and suggest some areas for further research.

PART I: OVERVIEW OF KANT'S ETHICAL THEORY:

GOOD WILL, DUTY, AND UNIVERSALIZABILITY

In this section, I'll outline a common and widely accepted interpretation of Kant's moral philosophy based on practical reason.[1]

According to Kant, imperatives are the formulas of the commanding principles of practical reason. Imperatives are expressed by an *ought* and state that some action possible to an agent would either be good to do or good to refrain from doing. Reasons for acting on any imperative derive either from pure reason or from inclinations (emotions, desires, etc.).

For Kant, there are two types of imperatives: hypothetical and categorical. Imperatives that are derived from inclination and are directed toward some subjective external purpose are called hypothetical imperatives. A hypothetical imperative states that an action is good for some possible or actual purpose. Hypothetical imperatives "represent the practical necessity of a possible action as a means for attaining something else that one wants (or may possibly want)" (414, pg. 25).

A categorical imperative, on the other hand, represents an action as objectively necessary in itself, without reference to any other purpose or end (414 and 415, pg 25). A categorical imperative commands certain actions with no condition—nor any other

---

[1] All references to Kant's work in this section, unless otherwise stated, refer to *Grounding for the Metaphysics of Morals* (1785/1993) translated by James W. Ellington, 3rd edition (Indianapolis, IN: Hackett Publishing Company). The first number in all citations refers to the page number in Volume IV of the Royal Prussian Academy edition of Kant's works, while the second number refers to the page number in Hackett's 3rd edition.

purpose to be attained by those actions—independently of one's subjective ends.[2] A

categorical imperative, then, prescribes actions that are intrinsically worthy – actions that

are what Kant refers to as 'ends in themselves.' Categorical imperatives command acts

that are objectively necessary in themselves, and their commands apply to us (human

persons) simply because we have rationality.[3]

This brings us to one of the most important parts of Kant's account of the nature

of morality – the intrinsic unconditional worth of a good will. According to Kant, there is

only one thing that is good under all conditions, and this a good will. No other thing, not

any of the virtues nor even happiness, is good under all conditions. Any of the virtues,

any talents or gifts of fortune, even happiness can be bad or wrong under some

conditions. All of these other conditionally 'good' things and all activities need a good

will for their good use. Only a good will, he says, "is good not because of what it effects

or accomplishes, nor because of its fitness to attain some proposed end; [only a good

will] is good only through its willing, i.e., is good in itself" and is esteemed higher than

anything that it could ever bring about to favor some contingent inclination, or even the

sum of all inclinations (happiness) (394, pg. 7). In fact, a good will itself is to be

esteemed above all else, even if it cannot accomplish its purpose:

> [I]f with the greatest effort it should yet achieve nothing, and only the good will
>
> should remain (not, to be sure, as a mere wish but as the summoning of all the
>
> means in our power), yet would it . . . still shine by its own light as something

---

[2] Also important here is the distinction between subjective ends ("which rest on
incentives") and objective ends ("which depend on motives valid for every rational
being") (427, pg. 35).
[3] According to Kant, both hypothetical and categorical imperatives apply only to rational
beings (425, pg. 33).

which has its full value in itself. Its usefulness or fruitlessness can neither

augment nor diminish this value. (394, pg. 8)

And it is only reason that can influence the will and produce a will that is good in itself;

nothing as subjective and contingent as inclinations can produce a will good in itself

(396, pg. 9).

All human beings possess (the possibility of acting from) a pure will, but also

possess desires and inclinations that conflict with it. Even a good will is prone to

'subjective restrictions or hindrances' (397, pg. 9). The good will along with its

hindrances are gathered together under the concept of duty. It is these 'subjective

restrictions and hindrances' that Kant uses to bring out what duty is. Kant asserts that

"duty has to be a practical, unconditioned necessity of action" that holds "for all rational

beings (to whom alone an imperative is at all applicable) and for this reason only can it

also be a law for all human wills" (425, pg. 33).

Kant asserts that there are four different categories of action in regards to duty:

there are (1) actions that are contrary to duty (that all rational people would consider

wrong, i.e., killing of innocents); (2) actions that are *in accord with* duty yet done from

some indirect (mediate) inclination (such as obeying the speed limit to avoid having to

pay a fine); (3) actions that are *in accord with* duty but done from some direct

(immediate) inclination (i.e., preserving one's life because one desires to not die); and

finally (4) actions that are done *from* duty, even when all inclinations are for something

else (397-398, pp. 10-11). Kant advocates this fourth category of action as being the test of a will's possible goodness.[4]

At this point, Kant outlines three propositions of morality: First, to have moral worth, an action must be done *from duty* (not from inclinations). This is primary to Kant's entire project. One's inclinations can change because inclinations are a contingent matter, and morality cannot be based upon something so contingent and subjective.

All actions that are done willingly, Kant says, are willed based upon some subjective principle of volition; these subjective principles of volition are called 'maxims' and take the form: I will do action A in order to achieve purpose P. Kant's second proposition of morality claims that an action that is performed from duty doesn't have its moral worth in the purpose to be attained by it (the content of the maxim), but in the (form of the) maxim in which it is determined. For example, someone might donate surplus food to a soup kitchen. However, whether this action has moral worth or not does not depend on the 'good' accomplished by the action (hungry people being fed); rather, whether the action has moral worth depends on whether the action is motivated from inclination or from duty. Thus, there is nothing specific in a list of possible actions or possible purposes that makes an action good or bad.

The third proposition, which Kant claims follows from the previous two, states that duty is the necessity of an action done out of respect for the moral law (398-400, pp. 11-13). Inclinations and the effects of actions are too subjective and contingent to command moral action. Only the moral law can command necessarily:

---

[4] It is important to note that acting from duty as Kant sees it rules out both motivations of selfishness *and* motivations of more benevolent feelings.

Thus the moral worth of an action does not lie in the effect expected from it nor in any principle of action that needs to borrow its motive from this expected effect. For all these effects … could have been brought about also through other causes and would not have required the will of a rational being, in which the highest and unconditioned good can alone be found. Therefore, the pre-eminent good which is called moral can consist in nothing but the representation of the law in itself, and such a representation can admittedly be found only in a rational being insofar as this representation, and not some expected effect, is the determining ground of the will. This good is already present in the person who acts according to this representation, and such good need not be awaited merely from the effect. (401, pp. 13-14)

But if the law does not determine the will through any effect or expected effect, what is the law that commands the will to moral action? Because Kant denies that the particular content of any maxim determines any action's moral worth, he claims that only universal conformity of the will's actions to the law as such (not to any law determining particular actions) can serve the will as a moral principle for acting. Therefore, duty is constituted by the necessity of acting from pure respect for the practical law, and every other motive is inferior to duty.

Now back to categorical imperatives: Only a categorical imperative can direct moral action. But Kant also concludes that only the universal conformity of the will's actions to the moral law as such can serve the will as a necessary principle. According to Kant, a categorical imperative contains both the objective principle of the law as such, and the subjective principle of acting (maxim) of a particular human being. Because

Kant holds that an end that is objective must hold for all rational beings (in other words, it must be *universalizable*), it is clear how Kant arrives at his primary formulation of the categorical imperative, the Formula of Universal Law:

> Act only according to that maxim whereby you can at the same time will that it should become a universal law. (421, pg. 30)

Here Kant is saying that the principle of moral law dictates that we should act only according to maxims that we can also will every other rational being to act on. This is the supreme principle of practical reason, according to Kant, and it provides guidance for our actions.

PART II: A KANTIAN ANALYSIS OF KNOWLEDGE

A: Kantian Theoretical Reason

Whereas practical reason deals with deliberations about action, theoretical reason can be understood as related to matters of fact and their explanation. Therefore we can associate practical reason with morality and theoretical reason with the realm of epistemology. Part of my project in this paper, however, will be to show how these two kinds of reason are related, and thus how Kant's moral theory might help us better answer some important questions about knowledge and its value.

In the *Critique of Pure Reason*, Kant criticizes the traditions in metaphysics that rely on the supersensory and the transcendent (that is, anything beyond what is revealed by the senses), and shows that reason by itself (without sensory and empirical experience) is prone to error. Kant's critique of metaphysics and rationalism leads to a view that the only knowledge we can really have is limited to what we can perceive through the senses (and perhaps evaluate through reason, but that which is a priori is not knowledge). Thus Kant denies reason's claims to transcendent insight because claims about God and the soul, etc., are exactly the kinds of claims that not all rational beings can accept. Because people cannot experience these things through their senses or empirically, Kant acknowledges that people will come up with variations of these ideas. The conflicting versions of people's ideas on these matters will cause people to fail to communicate with each other, fall into conflict, or submit to some despotic and unreasoned authority. Thus, reason by itself – without regard for empirical or sensory experience – is prone to err when attempting to reach the goal of knowledge.

Reason's role is regulative, not constitutive, Kant argues. That is, reason is regulative in that it is not by itself a source of knowledge, but it helps us to correct errors and guides us in obtaining further insight. So while Kant wants to delimit the bounds of reason, he doesn't want to argue that it has no role in our knowledge. For example, Kant also argues that, as regulator, reason is the arbiter of truth in all judgments. We perceive things through our senses all the time, but our senses do not judge. Kant observes that it is only when judgment enters the picture that we can be mistaken: "It is correctly said that the senses do not err; yet not because they always judge correctly, but because they do not judge at all" (A293). And reason is what helps us to understand whether our judgments are correct or not. For instance, suppose I dream that I inherited millions of dollars from an aunt who recently died. And suppose that I mistake the content of that dream as having actually occurred, and I believe that I really inherited millions of dollars from an aunt who recently died. But then I start to doubt my belief. To decide whether my belief is true or not, I need to reason and ask to what extent this belief connects up with my other beliefs and to the beliefs of other people. If my belief doesn't consistently connect up with my own other beliefs and with the beliefs of other people, then I have reason to think that my belief is wrong. As arbiter of judgments, reason's role is to seek unity and consistency in observations: as Kant says, "For the law of reason to seek unity is necessary, since without it we would have no reason, and without that, no coherent use of the understanding, and, lacking that, no sufficient mark of empirical truth…" (A651=B679). But more about this below.

While it isn't always apparent how Kant thinks that practical and theoretical reason relate to each other, in at least one section of the *Critique of Practical Reason*

entitled "On the primacy of pure practical reason in its connection with speculative reason" (5:119-121) he addresses this question. In the second Critique, Kant argues in general that pure practical reason has primacy over inclinations and desires; in this section he argues that it also has primacy over theoretical reason. He defines primacy as "the prerogative of the interest of one insofar as the interests of others is subordinated to it" (5:119). He claims "all interest is ultimately practical and even that of speculative reason is only conditional and is complete in practical use alone" (5:121). While exactly what Kant means in this section (and how it fits in with the rest of his philosophy) is open to interpretation, it seems from his arguments here that pure practical reason should guide (at least some of) our beliefs, as well as our actions.

There are three major concepts in Kant's practical philosophy that have use in a new Kantian analysis of knowledge: (i) Good Will, (ii) Duty, and (iii) Universalizability. I will show here how these three major concepts can be understood in a new light as pertaining not only to the realm of practical reason, but also to theoretical reason.

B. Good will: Following moral and epistemic duty.

Kant begins his *Groundwork for the Metaphysics of Morals* by establishing the unconditional and final value of a good will. A good will, he writes, is the only thing that can be regarded as good without qualification, and a good will has intrinsic unconditional worth (393 – 394). A good will is not good because of its effects or accomplishments, even if it produces good ends; a good will is good in itself, non-instrumentally. In other words, a good will might have instrumental value, but that is not its only or primary value. As Kant writes: "When [a good will] is considered in itself, then it is to be

esteemed very much higher than anything which it might ever bring about merely in order to favor some inclination, or even the sum total of all inclinations" (394).

While Kant goes on to explain how good will is related to moral duty, it is clear that good will is of significant value in its connection to reason, in general. Kant asserts that reason has a special role as it pertains to good will; that is, only reason can shape the will and produce a will that is good in itself. In fact, further support for understanding practical reason as guiding at least some of our beliefs can be found regarding good will. Practical reason is to guide all actions, and all activities need a good will for their good use. From my understanding of Kant's account of reason, he assumes doxastic voluntarism – that is, that, at least to some extent, we can decide what to believe and effectively will to change our beliefs. From such a stance, it is clear that will is important in forming or regulating at least some of our beliefs. Such doxastic voluntarism includes that rational beings have direct voluntary control over at least some of their beliefs, and that rational beings have indirect voluntary control over many of their beliefs by, for example, examining evidence, seeking justification, opening their beliefs to others' scrutiny, conducting research and improving methodologies, etc.

We can see that rational beings do have some direct form of voluntary control over the beliefs they form in light of sensory experiences. For example, take someone who is undergoing some sort of experimental pharmaceutical treatment for a long-term illness. The drugs she is taking are known to cause hallucinations, and she has been made aware of this possible affect. Perhaps she very strongly perceives that she is being

attacked by a horde of spiders, yet judges that she is not being attacked by a horde of spiders.[5]

It is incontrovertible that rational beings have indirect voluntary control over their beliefs, at least in the sense of being able to use their reason to examine evidence, seek justification, weigh reasons, subject their judgments to the scrutiny of others, etc. In this view, then, at least some beliefs are the type that can be acquired, guided and monitored by an intention or by duty – that is, by one's rational will. In this view, deliberating, seeking justification for belief, and all kinds of other epistemic processes are clearly activities through which we exercise intention. As such, according to Kant's account of practical reason, practical reason and good will must guide the good use of these activities. Epistemic good will involves using reason to guide and regulate our beliefs; this will include rational beings being able to recognize reasons as the right (or wrong) kind of reasons for belief. The right kind of reasons for belief might include reliable evidence, well-examined justification, and/or testimony of reliable authorities, for example. The wrong kind of reasons for belief will be those that are contingent on our inclinations, desires, and what we think will benefit us.

The aim of epistemic activities is one or more fundamental epistemic goods.[6] One possible fundamental epistemic good is knowledge. It is widely agreed that

---

[5] Such a person would not be blameworthy should she judge her strong sensory experience to indicate she was actually being attacked. However, if she uses her strength of will, along with the various information and understanding she has of the side effects of the drug, she is praiseworthy for following her epistemic duty in a very difficult situation.

[6] I'm assuming Duncan Pritchard's understanding of fundamental epistemic goods/goals here: for Pritchard, a fundamental epistemic good is any epistemic good whose specifically epistemic value is *at least sometimes* not simply instrumental to any other *epistemic* good. Non-fundamental epistemic goods are always instrumentally valuable

knowledge requires at least true belief, whatever else it might include. But according to the principles laid out above (i.e., doxastic voluntarism and that all activities need a good will for their good use), knowledge as a fundamental epistemic goal complying with epistemic duty must also always include a good will (epistemic good will as explained above). Such a Kantian account of knowledge might include other elements, but it seems clear that one cannot be complying with epistemic duty without a good will. It is important to note that a rational being might have a good will without knowledge; this is the case because knowledge also involves (at least) true belief, and an agent might have good will without true belief. In this analysis, however, as opposed to some other epistemic accounts, true belief would not be the singular fundamental epistemic good.

To understand why epistemic duty includes a good will, I turn to Kant's explanation of actions that merely are in accord with duty and those that are from duty. True belief is the epistemic counterpart to action merely in accord with duty. That is, true belief might be attained for a variety of reasons and motives that may or may not involve good will. For instance, someone might hold a true belief about the existence of God, but only hold that belief because it makes her life more comfortable, or only because she passively and lazily assents to the supposed epistemic authority of another without subjecting those judgments to the scrutiny of reason. Knowledge, on the other hand, is the epistemic counterpart to action from duty. While knowledge includes a true belief, that true belief must be held through good will for it to count for knowledge. If I have knowledge, I don't hold the true belief merely because I'm inclined to, or because it lines up with my pursued ends. I might have true beliefs for such reasons, and that is

relative to some other epistemic good. Both fundamental and non-fundamental epistemic goods might still be instrumentally valuable relative to some further non-epistemic good.

instrumentally valuable. But for knowledge, my belief must be true, and it must be held because I'm complying with my epistemic duty.

C. Epistemic duty: Universalizability

But what exactly does epistemic duty consist in, and what does a good will look like when it comes to epistemic goals? We know from Kant's moral philosophy that the Categorical Imperative is the supreme principle of practical reason. But as Onora O'Neill[7] has pointed out, if practical reason has primacy over theoretical reason, and the Categorical Imperative is the supreme principle of practical reason, then the Categorical Imperative is the supreme principle of reason, full stop. Kant doesn't explicitly say this, but in a footnote to his 1786 essay "What is it to Orient Oneself in Thinking?" he writes:

> To make use of one's own reason means no more than to ask oneself, whenever one is supposed to assume something, whether one could find it feasible to make the ground or the rule on which one assumes it into a universal principle for the use of reason. (8:146n)

This clearly parallels the first formulation of the Categorical Imperative (The Formula of Universal Law): "Act only according to that maxim whereby you can at the same time will that it should become a universal law" (421, pg. 30). Kant seems to be saying: ***believe (or think) only according to that maxim that could be a universal law***. Where Kant sees that it is categorically imperative to will only those actions that rational beings can will (i.e., that can be universalized), the epistemic counterpart might be that we should only believe those things that rational beings can/should believe (i.e., that can be

---

[7] Onora O'Neill, *Constructions of Reason*, Cambridge: Cambridge University Press, 1989.

universally believed) on the grounds available to us. Thus, I know p on grounds G only if any rational being would/should believe p on grounds G. If a rational being would refuse to believe p on grounds G, then even if I believe p on grounds G, I don't really *know* p. I might be merely entitled to believe p on grounds G (e.g. believe that God exists based upon available evidence or personal experiences that I have), but am not rationally required to believe p on these grounds; it's optional, and hence not knowledge.

It is interesting to note here that, according to this understanding of epistemic duty our direct voluntary control over belief seems to be complying with epistemic duty best when it has more of a regulative/negative character than a positive/constitutive one. That is, we are best following epistemic duty when exercise more direct control over what *not* to believe than over what to believe. The control we have over what to positively believe seems to be more indirect – in the sense of seeking justification, examining evidence, evaluating methodologies, and other such processes. Our reason judges what reasons we have to believe P, to believe not P, or to not believe P. If through our epistemic good will we judge that there are more or better reasons (of the right kind) to believe P than there are to believe not P, then we are permitted—and perhaps obliged—to believe P. If in the same way we judge that there are more or better reasons (of the right kind) to believe not P, we are permitted—and perhaps obliged—to believe not P. When we are permitted to believe P (or to believe not P), we are also permitted to not believe P, because mere permission entails that there is not strong empirical evidence for P (or not P), and thus we are not required to believe P (or not P, whatever the case may be).

Beliefs about the trans-sensory—for example, beliefs about God—are an interesting case here. According to Kant's metaphysics and account of theoretical

knowledge, we know that there is not enough empirical evidence about God's existence

for us to have knowledge about it. Thus, we cannot be required to believe that God does

exist (P) or that God doesn't exist (not P). We're only responsible for the belief about

which there is some empirical evidence available. What we can be held accountable for

is something like failing to conduct an adequate investigation into or failing to be open to

certain propositions because of certain (conscious or unconscious) inclinations like

laziness or fear. This is important in that, when we are following our epistemic duty and

exercising epistemic good will, we don't will to believe certain things (P or not P) based

upon what is at stake for us. We have the capacity to do this—to exercise direct

voluntary control over our beliefs; but we mustn't believe things simply because doing so

would fulfill our contingent desires and inclinations. Rather, we use our reason and

epistemic good will to avoid such maxims that are not from epistemic duty.

Epistemic duty involves some very specific things in this analysis. O'Neill (1989

and 1992) discusses three "maxims of common human understanding" (reason) that

appear in the Critique of Judgment[8] that seem to be closely related to the Categorical

Imperative and support the claim that human knowledge is subject to universalizability.

These three maxims are: (i) to think for oneself, (ii) to put ourselves in thought in the

place of everyone else, and (iii) to always think consistently with oneself (§40).

According to Kant, the first maxim protects against prejudice—that is, it is a maxim

against the passivity of reason. The passivity of reason can cause us to be blindly guided

by others, who may or may not have good will or well-reasoned judgments. The second

maxim reminds us that while we should not blindly follow those who claim to have

---

[8] But they also appear much earlier in Kant's *Lectures on Logic* ("The Jäsche Logic," 57) and also are referenced in "What is Enlightenment?"

authority, we must still follow principles of thinking and judgment that are open to other rational beings. The second maxim also helps us to enlarge our thought. Kant writes that it marks a person of enlarged thought "if he disregards the subjective private conditions of his own judgment, by which so many others are confined, and reflects upon it from a *universal standpoint* (which he can only determine by placing himself at the standpoint of others)" (§40). The third maxim ensures consistency in our thought, but Kant writes that it is the most difficult of the three to attain, and can only be achieved by the combination of the first two maxims and after regular observance of them has made our adherence to them automatic. We can achieve consistency in thought and judgment only when we genuinely try to judge for ourselves and subject our judgments to the examination of others. However, all three of these maxims work together to give content to epistemic duty; any one in itself, or in combination with another but without the third, is insufficient to guide our beliefs in compliance with epistemic duty.

These three maxims show the implications of the Categorical Imperative as the supreme principle of reason. As O'Neill writes, reason is "the lawlike guidance" of both thinking and doing[9]; this supports the claim to the unity in structure of practical and theoretical reason. The maxims also flesh out what our epistemic duty consists in—that is, what a good will requires in regards to epistemic goals and regulating our epistemic activities.

---

[9] Onora O'Neill, "Vindicating Reason," in *The Cambridge Companion to Kant*, edited by P. Guyer, Cambridge: Cambridge University Press, 301.

PART III: THE VALUE OF KNOWLEDGE

A. Outlining the problems…and one promising response to the problems

One long-standing project in epistemology is the attempt to understand what knowledge is and why it is distinctly valuable. Above, I presented an account of what knowledge is based on Kant's account of Practical and Theoretical Reason. In this section, I will deal with the value of knowledge. First, I will explain what is meant by the claim that knowledge is distinctively valuable, outlining three specific challenges to the claim. Duncan Pritchard identifies these three challenges as (1) the Primary Value Problem, (2) the Secondary Value Problem, and (3) the Tertiary Value Problem; he also points out a related problem, (4) the Swamping Problem. I draw heavily on Pritchard's work on epistemic value in this section.

After addressing the different issues at stake in the claim that knowledge is distinctively valuable, I will outline Pritchard's argument for a virtue epistemic account that addresses these problems. I will show why his account doesn't come close enough to sufficiently addressing these problems, and argue why my account of knowledge is better than his as it pertains to these key Value of Knowledge questions.

Value of Knowledge Problem: Why Is Knowledge Distinctly Valuable?

Each of the different value problems addresses different readings of the claim "knowledge is distinctively valuable" (what I'll refer to as Claim KV). The Primary Value Problem deals with the interpretation of KV as "knowledge is more valuable than mere true belief." This version of KV is primary in that, if we are not able to explain why knowledge is more valuable than mere true belief, then we cannot address any of the

different versions of the value problem that arise from KV. One answer to this problem (suggested by Socrates) is that the value of knowledge over mere true belief is a greater practical value: that is, in general, one is more likely to achieve one's goals with knowledge than with mere true belief.

However, even if we take this as a successful answer to the Primary Value Problem, there is still more to the claim that knowledge is *distinctly* valuable. The Secondary Value Problem deals with the reading of KV as "knowledge is more valuable than any that which falls short of knowledge." The problem here is that, even if we can address the Primary Value Problem by explaining the greater practical value of knowledge over mere true belief (perhaps by adding some characteristic that increases reliability, such as *justification*), this doesn't show why we should distinctly value knowledge over that which falls short of knowledge – namely, mere true belief plus some added value-conferring criterion.

Furthermore, claiming that knowledge is *distinctively* valuable seems to suggest that "knowledge is more valuable than any epistemic standing that falls short of knowledge" not merely by degree, but by kind. The Tertiary Value Problem deals with this reading of KV. Pritchard claims that we often treat knowledge as being *precious*, unlike that which falls short of knowledge, in that knowledge is at least sometimes valuable for its own sake. He writes that this is a non-instrumental, and therefore *final*, value. If Pritchard is correct, then any adequate response to the Tertiary Value Problem must show that knowledge has final value—that is, value that is not merely instrumental to another end.

While many theorists have approached the value problem by starting out trying to answer the Primary Value Problem, Pritchard recommends a different approach. He recognizes that, if we can successfully answer the Tertiary Value Problem, then we will have sufficiently addressed the Primary and Secondary Problems as well. According to his approach – which I adopt below, our goal is to eventually find a way to show that knowledge is of final value in a way in which any epistemic standing which falls short of knowledge is not. But first, I'll spend some space here addressing what is known as "the Swamping Problem."

The Swamping Problem is similar to the Primary Value Problem, but is not merely a sub-reading of KV. The Swamping Problem involves a general axiological claim that, if a property (such as justification or a reliable process of production) is only instrumentally valuable relative to some other good (such as true belief), and if this further good is already present, no additional value is added by the presence of the merely instrumentally valuable property. The Swamping Problem for epistemic value applies to any epistemic proposal that treats some epistemic standing as instrumentally valuable to the good of true belief. According to Pritchard, the argument that 'whatever epistemic standing is in question has a practical value that mere true belief doesn't' won't be an answer to the Swamping Problem at all. The Swamping Problem is concerned not just with any kind of value, but specifically with epistemic value, Prichard argues; the main concern of the Swamping Problem is "how to make sense of the idea that knowledge is more valuable than mere true belief because it involves an epistemic standing which better serves our specifically *epistemic* goals – in particular, the epistemic goal of true belief" (8).

The Swamping Problem for epistemic value seems to arise from the contradictions among three highly intuitive claims:

(1) The claim that *true belief* is the sole fundamental epistemic good ("The epistemic value conferred on a belief by that belief having an epistemic property is instrumental epistemic value relative to the further epistemic good of true belief" (p. 12)); Pritchard refers to this view as "epistemic value T-monism" (p. 11);

(2) The claim about value that "If the value of X is only instrumental value relative to a further good and that good is already present, then it can confer no additional value" (p. 12); and

(3) The claim that knowledge is at least sometimes more epistemically valuable than mere true belief.

It is important to note here that Pritchard draws a distinction between *fundamental epistemic goods* and *final value*. According to Pritchard, a fundamental epistemic good is any epistemic good whose specifically epistemic value is *at least sometimes* not simply instrumental to any other *epistemic* good. Non-fundamental epistemic goods are always instrumentally valuable relative to some other epistemic good. Both fundamental and non-fundamental epistemic goods might still be instrumentally valuable relative to some further non-epistemic good.

The claim that true belief is the sole fundamental epistemic good means that all other epistemic goods are instrumental to the epistemic good of true belief. If we accept (1) and (2) as true, then necessarily we must reject (3) as false. The triad is inconsistent, and the Swamping Problem asks us to try to make it consistent.

Pritchard finds that there is no good reason to reject (2), and thus takes it for granted. We might deny (3), but we would need to deny it as part of a larger explanation of why we tend to think that (3) is true. This is the tactic of proponents of the 'practical response' to the Swamping Problem: they claim that the seeming greater epistemic value of knowledge over true belief is really just that knowledge is generally more practically valuable than true belief. To deny the very weak claim of (3), however, is to advocate a very strong alternative claim – that knowledge is never of greater epistemic value than mere true belief. The further problem with denying (3) is that (3) could claim that any epistemic standing is at least sometimes more epistemically valuable than mere true belief, which would contradict with (1); thus our denial of (3) entails our affirmation that no epistemic standing is ever of more epistemic value than mere true belief. This is a very strong thesis indeed. Even if a proponent of this strategy is willing to accept such a strong thesis, she must still (a) offer a plausible story as to why we might be wrongly inclined to think of (3) as true, and (b) argue that the Tertiary Value Problem is really not a problem at all (since it asks why knowledge is more valuable in kind – why it uniquely has final value – than that epistemic standings that fall short of knowledge, and we must deny the assumption that knowledge has final value since it is inconsistent with (1)). Pritchard does not deny this as a plausible strategy outright; instead he opts to examine whether alternative proposals to solving the inconsistent triad of the Swamping Problem are plausible *and* able to solve the Tertiary Value Problem without denying it as a problem.

If we accept (2) as given, and we're looking for alternatives to denying (3), then we must examine whether we can plausibly deny (1). We can deny (1) in two very

different ways: first, we can simply deny that true belief is the sole fundamental epistemic good without denying some sort of monism (that there is only one fundamental epistemic good), or we can reject epistemic value monism entirely and instead endorse epistemic pluralism (p. 16). Both of these views can be used to support the claim that knowledge is a fundamental epistemic good – either the only fundamental epistemic good (value monism) or one of multiple epistemic goods (value pluralism), one of which could also be true belief – and thus bypass the Swamping Problem. Pritchard asserts that whether we try to adopt value K monism, or we try to adopt value pluralism and argue that both knowledge and true belief are fundamental epistemic goods, we still need to demonstrate that knowledge is also finally valuable in order to be able to extend the pluralist response to the Tertiary Value Problem.

Does Knowledge Have Final Value? (Can we answer the Tertiary Value Problem?)
To address the Tertiary Value Problem (the problem of explaining why knowledge is more valuable in kind than that which falls short of knowledge), Pritchard argues, we are required to explain why knowledge has final value, unlike those epistemic standings that fall short of knowledge. Pritchard examines one account of why knowledge has final value, but concludes that it is ultimately unsuccessful at answering the value problem. However, the examination of this account brings to light some important concerns regarding epistemic value and defining knowledge.

The account that Pritchard finds most promising to explain why knowledge has final value is what he calls *robust* virtue epistemology—a view he attributes to Ernest Sosa, Linda Zagzebski, and John Greco. According to Pritchard, a virtue-theoretic

epistemology proposal is robust if "it attempts to exclusively analyse knowledge in terms of a true belief that is the product of epistemically virtuous belief-forming processes" (24). Pritchard focuses on three major benefits he attributes to robust virtue-theoretic accounts:

> (1) virtue-theoretic accounts line up with our strong intuition that knowledge is the product of an agent's reliable cognitive abilities,
>
> (2) robust virtue epistemology accounts seem to have the resources to deal with Gettier-style cases of knowledge-undermining epistemic luck (intervening luck, in this case), and
>
> (3) robust virtue epistemology accounts also line up with the common view that knowledge is a kind of cognitive achievement, in that achievements in general are successes because of one's ability, and virtue epistemology seems to be claiming the epistemic analogue.

Regarding (3), in this view knowledge is cognitive success because of one's cognitive abilities. Knowledge is then just a specifically cognitive type of achievement, but only in that one's success if primarily creditable to her abilities (rather than to luck or some other external factor that is not dependent on her abilities). This is the achievement thesis, but it is formulated here specifically in regards to knowledge. The achievement thesis combined with robust virtue epistemology entails the *knowledge-as-achievement* thesis (K=A), that knowledge is a type of achievement. This is significant because achievements are considered distinctively valuable – specifically, the successes that are described as achievements are of final value (not valuable merely in relation to any

further goal or practical benefit). This is the Value of Achievement thesis. Thus, Pritchard reasons, if we can show that knowledge is a type of achievement – unlike that which falls short of knowledge – then we may be able to show how knowledge has a distinct kind of value that lesser epistemic states lack. Pritchard outlines the argument thus:

> (P1) Achievements are successes that are because of ability. (Achievement thesis)
>
> (P2) Knowledge is a cognitive success that is because of cognitive ability.
>
> (Robust Virtue Epistemology)
>
> (C1) So, [from P1 and P2 we conclude that] knowledge = cognitive achievement.
>
> (K=A thesis)
>
> (P3) Achievements are finally valuable. (Value of Achievement thesis)
>
> (C2) So, [from C1 and P3 we conclude that] knowledge has final value.
>
> (Pritchard, 30)

While the inferences in this argument are valid, Pritchard argues that the robust virtue epistemological account of knowledge (P2) faces severe problems, particularly in its generation of the false K=A thesis (C1).

The K=A thesis is highly problematic. For example, there are cases in which knowledge doesn't involve the corresponding cognitive achievement, and likewise there are cases of cognitive achievements that don't equal knowledge (33). In particular, this account has problems with cases that involve another type of knowledge-undermining epistemic luck—environmental epistemic luck. This is the kind of luck involved in the much-discussed barn cases. In these types of cases involving environmental luck, an epistemic agent uses a reliable process to form a true belief and thus has a cognitive achievement; but nevertheless we don't count this cognitive achievement as knowledge

because the environment is epistemically inhospitable to the extent that the agent's belief is only true as a matter of luck. So while robust virtue epistemology can deal successfully with Gettier-style cases (intervening epistemic luck), it is not able to answer the problem of environmental epistemic luck. Thus, there is at least sometimes more to knowledge than mere cognitive achievement, which denies (C1). If (C1) is false, then either (P1) or (P2) must also be false. Pritchard concludes that (P2) must be false because the environmental epistemic luck challenge not only undermines the K=A thesis (C1), but also the specifically robust virtue epistemology claim that knowledge is cognitive success that is because of an agent's cognitive ability (P2).

Pritchard thus denies that exhibiting cognitive achievement is sufficient for possessing the corresponding knowledge. But he also denies that obtaining knowledge is sufficient to have a cognitive achievement. He introduces the 'Jenny' case – a case of testimony – in which young Jenny asks an adult passer-by for directions, and the adult (who has first-hand knowledge of the city) gives her the directions she has requested. It seems in this case that Jenny has knowledge – the correct directions that the passer-by gave her, but he claims that this acquisition of knowledge isn't *primarily* creditable to Jenny or to her cognitive agency. Since according to our account of robust virtue epistemology a true belief needs to be primarily creditable to the agent to constitute a cognitive achievement, Jenny may have knowledge in this case, but Pritchard thinks that it does not seem to exhibit a cognitive achievement. It seems that Pritchard might think that some cases of testimony (in which knowledge is gained) exhibit a type of cognitive achievement, but not in a strong sense. This, like the epistemic luck problem, is a challenge to both (C1) and to (P2). This is another way in which we see that knowledge

can't be identified with cognitive achievement. Thus not only is the K=A thesis

problematic (C1), but the robust virtue epistemological account of knowledge (P2) is also

significantly flawed. Even if (P2) weren't seriously flawed, it still wouldn't maintain its

appeal since without the K=A thesis it loses its unique resources to argue for the final

value of knowledge; the robust virtue epistemology account, then, seems unable to

address the Tertiary Value Problem that we set out to solve. Pritchard concludes that,

since robust virtue epistemology was the only promising view for accounting for why

knowledge is distinctly and finally valuable, but it proved unsuccessful in developing

such an account, the Tertiary Value Problem cannot be overcome.

Pritchard's Solution: Knowledge is Sometimes *of Final Value*

With these new insights from examining robust virtue epistemology, Pritchard develops

what he calls an anti-luck virtue epistemology theory that he claims incorporates the

strengths and major intuitions of robust virtue epistemology, but also avoids its

weaknesses and problems. Since Pritchard sees the robust virtue epistemological account

as most promising to solve the Tertiary Value Problem, Pritchard concludes that we can't

sufficiently answer why knowledge is more valuable in kind than that which falls short of

knowledge. Therefore, Pritchard does not try to show how this new anti-luck account can

address the Tertiary Value Problem, but he does aim to show where the intuition that

grounds the Tertiary Value Problem – the intuition that knowledge is distinctively

valuable – comes from.

Pritchard argues that, while robust virtue epistemology and other virtue-theories

come close to offering a correct understanding of knowledge, ultimately they fail in not

recognizing that there are two master intuitions about knowledge. Other virtue accounts

accommodate for the ability intuition (that knowledge is the product of an agent's

cognitive abilities when knowledge as cognitive success can be significantly credited to

the agent). But only anti-luck virtue epistemology also accommodates the anti-luck

intuition – that if one has knowledge then it is not the case that one's true belief could

easily have been false (52). These two master intuitions impose distinct demands on

one's theory of knowledge. Anti-luck virtue epistemology gives equal weight to both of

these fundamental intuitions, and incorporates both an anti-luck and an ability condition

that constrain what is considered knowledge.

The account of knowledge offered by anti-luck virtue epistemology is structured

as follows: knowledge is true belief that couldn't have easily been false that arises out of

a reliable cognitive character, such that one's cognitive success is **to a significant degree**

creditable to one's cognitive character (55). The 'true belief that couldn't have easily

been false' element is the anti-luck condition in the view, and virtue-theoretic element is

the ability condition. This formulation of the ability condition is different than the ability

constraints of the robust virtue epistemology account he introduced earlier; this version

does not demand that the cognitive success be entirely because of an agent's cognitive

ability, but rather only that the cognitive success should be *to a significant degree*

creditable to one's cognitive character. Thus anti-luck virtue epistemology shies away

from a strictly achievement account of knowledge.

The ability condition formulated in this way accommodates 'Jenny' type cases in

which an agent's cognitive success isn't strictly creditable primarily to her cognitive

abilities. As long as Jenny's (or another agent's) true belief is at least partly creditable to

her, we can count her as having knowledge. The degree of cognitive ability required in order to know is dependent on how epistemically friendly the environment is. If Jenny is in an environment where the anti-luck condition can be easily satisfied (for instance, a place in which there are many honest and reliable informants), then the bar for her cognitive ability in order to be able to know is quite low. But if Jenny is in an epistemically unfriendly environment (perhaps where there are lots of deceptive informants), the degree of cognitive ability required for her to gain knowledge is quite a lot higher.

Anti-luck virtue epistemology doesn't only accommodate 'Jenny' type cases, but it also responds well to various problems within epistemology, such as in response to Gettier-style intervening epistemic luck cases, and the even more problematic cases of environmental epistemic luck. Pritchard also notes that this account with its anti-luck condition can deal with 'lottery' cases, in which an agent's true beliefs based upon their cognitive abilities could very easily have been false. Since anti-luck virtue epistemology with its two major constraints of knowledge – the anti-luck condition and the ability condition – seems to successfully respond to a wide variety of problems in epistemology, it seems that there is much initial support for the view.

But Pritchard still must address how anti-luck virtue epistemology can handle the value problem. Pritchard concluded in his discussion of robust virtue epistemology that the proposal most likely to be able to account for the final value of knowledge (and thus adequately respond to the Tertiary Value Problem) is actually untenable; thus, he claims that anti-luck virtue epistemology need only respond to the Primary and Secondary Value Problems and be able to diagnose why the Tertiary Value Problem appears as a problem.

Pritchard thinks that anti-luck virtue epistemology has very good resources to be able to offer a diagnostic account of the Tertiary Problem, and to also respond to the Primary and Secondary Problems. To address the Tertiary Problem, Pritchard reminds the reader that in the previous chapter he declared that knowledge is not finally valuable, but that cognitive achievements are distinctively (finally) valuable. Because knowledge sometimes (but not always) involves cognitive achievement, then, knowledge is also *sometimes* finally valuable when it is a cognitive achievement. Pritchard notes that the exemplary cases of knowledge we think about tend to be cases in which the knower has the corresponding cognitive achievement (for example, we don't think of 'Jenny' type testimony cases as being paradigm cases of knowledge, he says). Because knowledge and cognitive achievement often overlap, it is natural to suppose that knowledge itself is of final value, even though it is actually cognitive achievement that has final value. Because fundamental epistemic goods are those which are at least sometimes valuable, knowledge can still be a fundamental epistemic good under this account. So in this sense, knowledge sometimes has final value – when it is also a cognitive achievement.

Pritchard claims that Craig's[10] account of the genealogy of the concept of knowledge suggests the beginnings of an adequate answer to the Secondary Value Problem, and thus also the Primary Value Problem. In this account, knowledge has a distinctive epistemic standing that is particularly instrumentally valuable. Thus we could argue that knowledge is of more value instrumentally than that which falls short of knowledge. This account should be developed further.

---

[10] Edward Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford University Press, Oxford, 1991.

Finally, Pritchard addresses how anti-luck virtue epistemology responds to the Swamping Problem. As I mentioned above, knowledge is sometimes of final value (because it sometimes constitutes a cognitive achievement), and is therefore a fundamental epistemic good. Thus, according to Pritchard, it doesn't matter that knowledge is not always of final value; knowledge at least sometimes has additional value over mere true belief, which does not have final value. Pritchard thinks this a good answer to the Swamping Problem.

B. Why Pritchard's response is inadequate

Pritchard concludes that the Tertiary Value Problem (the problem of explaining why knowledge is more valuable in kind than that which falls short of knowledge) is unsolvable because, in his account, knowledge doesn't always have final value. While he thinks the Tertiary Value Problem unsolvable, Pritchard thinks that his account adequately addresses the Secondary Value Problem (the problem of why knowledge is of more epistemic value than any proper subset of its parts) and the Primary Value Problem (the problem of why knowledge is of more epistemic value than true belief) by showing that knowledge is an exercise of skill. Because skill and ability are sometimes cases of achievements, and achievements have final value, he says that knowledge sometimes has final value. If knowledge sometimes has final value, then sometimes it is more valuable than any proper subset of its parts and more valuable than true belief, Pritchard argues.

The problem with Pritchard's answer is that he seems to want to solve the Tertiary Value Problem for all cases, which is why he thinks the 'knowledge is sometimes of final value' answer fails to adequately address this problem. But he seems satisfied to solve

the Primary and Secondary Value Problems for only some cases. He needs to explain more explicitly why he thinks this is an adequate response for the Primary and Secondary problems, but not the Tertiary Problem.

It would seem that an answer to the Primary Value Problem should explain why knowledge is always more valuable than true belief. Likewise, it seems that an answer to the Secondary Value Problem should explain why knowledge is always more valuable than any proper subset of its parts. If there is an account of knowledge that answers both the Primary and Secondary Value Problems for all cases, it would seem to be a better account of the value of knowledge. If another account of knowledge could explain why knowledge always has final value (and thus adequately answer all three of the value problems, along with the Swamping Problem), this account would be clearly superior to Pritchard's as far as addressing these key Value of Knowledge problems.

My Kantian account of knowledge outlined in Part II above offers an response for all three of the value problems, solving the Tertiary and Primary Problems along with the Swamping Problem not in just some but all cases, and offering a promising start to addressing the Secondary Problem. My account is thus superior to Pritchard's for addressing the key questions of the Value of Knowledge. I will outline how my account better addresses these problems in the following section.

C. *How my Kantian analysis of knowledge better addresses these problems*

As stated above, Pritchard recognizes that, if we can successfully answer the Tertiary Value Problem, then we will have sufficiently addressed the Primary and Secondary Problems as well. To accomplish this, he says our goal is to show that knowledge is of

final value in a way in which that which falls short of knowledge is not. Pritchard could not achieve this goal with his anti-luck virtue epistemology. But my Kantian analysis of knowledge shows that knowledge is clearly of final value in a way that true belief (or any other epistemic state that falls short of knowledge) is not.

According to my analysis of knowledge, knowledge always includes both good will and true belief. Because good will has final value, and knowledge must always include an agent utilizing good will, knowledge always has final value. True belief, on the other hand, does not have final value. True belief is of conditional value—unlike knowledge—and is (at least sometimes) merely instrumental to knowledge as a rational epistemic agent can hold a true belief but lack good will. As I argued above, true belief is the epistemic counterpart to Kant's category of action merely 'in accord with duty' because true belief might be attained for a variety of reasons and motives that may or may not involve good will. Knowledge, on the other hand, is the epistemic counterpart to Kant's category of action 'from duty.' While knowledge includes a true belief, that true belief must be subjected to the universalizability criterion of the epistemic counterpart of the Categorical Imperative, along with the other maxims of 'common human understanding' examined above. For it to count that I have knowledge, my belief must be true, and it must be held because I'm complying with my epistemic duty.

Because my analysis of knowledge shows why knowledge always has final value while true belief does not, it adequately answers two of the three key value problems for knowledge, along with the Swamping Problem:

Primary Problem: Because knowledge has final value (and true belief does not), it is more valuable than mere true belief;

Tertiary Problem: Because knowledge has final value (and true belief and other epistemic standings do not), knowledge is more valuable *in kind* than any epistemic state that falls short of knowledge.

Swamping Problem: Because knowledge has final value it is at least sometimes not simply instrumental to any other epistemic good. This makes knowledge a fundamental epistemic good. Whether or not there are other fundamental epistemic goods that include true belief (that is, whether or not my Kantian account of knowledge turns out to be a monist or pluralist account of fundamental epistemic value), it is clear that knowledge is not merely instrumental to true belief or any other epistemic good. This addresses the epistemic version of the Swamping Problem.

In my account, knowledge is a fundamental epistemic good. I find it likely that there are other fundamental epistemic goods, including true belief and possibly some other epistemic states. The important thing for addressing the value of knowledge and epistemic Swamping Problem is that true belief is not the only fundamental epistemic good under this account. Knowledge is of final value, and is thus not merely instrumental to any other epistemic good, including true belief.

While Pritchard thinks that successfully addressing the Tertiary Value Problem will also give an adequate response to both the Secondary and Primary Problems, I have some doubt.

The Secondary Value Problem deals with the reading of KV as "knowledge is more valuable than any that which falls short of knowledge." The problem here is that, even if we can address the Tertiary Value Problem by explaining the greater value of knowledge over mere true belief by adding some characteristic such as good will, this

might not sufficiently show why we should distinctly value knowledge over that which falls short of knowledge – namely, mere true belief plus some added value-conferring criterion. The Secondary Problem as Pritchard lays it out leaves me perplexed because it is not clear whether this reading of KV means that we must show why knowledge is more valuable than any epistemic standing that falls short of knowledge, or whether we must show why knowledge has more value than both any epistemic standing that falls short of knowledge (such as true belief) and any other proper subset of its parts (including why knowledge is more valuable than good will + true belief).

In either reading, I think it is important in solving the Secondary Value Problem to address the anti-luck quality that we seem to think that knowledge has. Similar to the virtue epistemology account that Pritchard criticizes and the one he offers, my Kantian analysis of knowledge is open to a credit view. That is, Kant thinks that acting from duty is morally credit-worthy for the agent, but acting merely in accord with duty is not credit-worthy (not blameworthy necessarily, either). The epistemic counterpart here would be that true belief attained by complying with epistemic duty (that is, through good will) is of credit to the rational being.

Let's lay this out a bit more clearly. There seem to be three different types of luck cases that we need to address: good luck cases of true belief, bad luck cases of false belief, and good luck cases of true belief that could have easily been false. Good luck cases of true belief are not of credit to the agent. For example, suppose Peter looks at his friend and says (and believes that), "There are four one-dollar bills in your left rear pocket." Let's suppose that Peter doesn't have any good reason for thinking there are four one-dollar bills in his friend's left rear pocket—he's just guessing. But luckily

enough, Peter has made a good guess; there are indeed four one-dollar bills in his friend's specified pocket. Even though Peter has obtained a true belief, it is not of credit to him in any significant way. He hasn't complied with his epistemic duty in any meaningful way, and though we don't find his guess necessarily blameworthy, it is obviously not credit-worthy. He hasn't complied with his epistemic duty in any significant way, thus his lucky true belief is not of credit to him, and he does not have knowledge.

There are other cases, however, in which a rational being might be complying with her epistemic duty—which is credit-worthy, but she might still not have knowledge because she lacks true belief. One example of this kind of bad luck case with good will might be a type of case in which an agent uses resources available to her to examine the evidence of a particular proposition, but some sort of knowledge-undermining epistemic luck intervenes in an unpredictable way to provide her with a false belief. For instance, let's suppose that Maggie sees what appears to be a class of milk on the kitchen table. She doesn't have any strong reason to think it's not a glass of milk, but doesn't just assume that it is a glass of milk, and goes over to look more closely at it, smell it, etc. After examining the liquid in the glass and noticing that it conforms to several of the most obvious characteristics of milk (color, consistency, smell), she forms the belief, "There is a glass of milk on the table." She has complied with her epistemic duty by examining the empirical evidence available to her, etc.; however, the glass is actually filled with some non-milk liquid that normally would not be in her kitchen but bears a striking resemblance to milk. In this situation, Maggie is credit-worthy for complying with her epistemic duty (she sought justification), but she doesn't have knowledge because she lacks true belief.

But these two types of luck cases are easy to deal with. The difficult epistemic luck cases are those such as good luck cases of true belief that could easily have been false. This includes both Gettier-style cases and barn cases. My account doesn't at this point offer a specific anti-Gettier (or anti-barn case) component, though it is perhaps open to further development. Thus, Gettier-style cases and barn cases might still present questions for my account, but these kinds of cases don't seem to me to be viciously problematic for my account. However, epistemic duty might require that for situations in which it is very important for other reasons to be right about belief, that you investigate as much as you can. For instance, it is very important (i.e., for the sake of the health and life of a child) for an agent to know that she's not giving spoiled milk to a baby. In this scenario, epistemic duty might require more than it otherwise would in service to some other requirement—the moral requirement to do no harm, perhaps, or to not be negligent. Her epistemic duty requires in this situation that the epistemic agent investigates as much as she can. For an average agent (someone who doesn't have special milk-examining skills), this would include not just looking at the milk (even though we think of sight as a fairly reliable sense), but also using the other senses to smell, taste, touch, etc., the milk to make sure that it is not spoiled.

However, there could be some situations in which, though we come to the correct judgment, we still might easily have been wrong. If the true belief is obtained by the epistemic agent because of her compliance with good will (that is, because she sought justification, examined the evidence, followed the maxims of common human understanding as they applied to the situation, etc.), she is still credit-worthy and has gained knowledge, even if she might have easily been wrong. This seems to me a

sufficient response to the Gettier and barn cases, even if it does not solve them in every situation, since the problems such cases involve are not always vicious.

Even if this is not a sufficient response to the luck cases, I have sufficiently addressed the Secondary Problem of showing why knowledge has not just final value because it includes a good will, but also has more value than both true belief and any other proper subset of its parts (including good will + true belief). Knowledge has value above and beyond any subset of its parts because of the anti-luck quality that arises from its components. The anti-luck quality of my view is not ad hoc; it arises out of the components of knowledge, but isn't a separate component from them.

CONCLUDING REMARKS

In this paper, I have presented a novel analysis of knowledge, drawing from Kant's account of practical and theoretical reason. While this approach may seem similar to some virtue epistemic accounts, it is unique in its explicit relying on duty and the Kantian universalizabilty criterion. This Kantian account of knowledge better addresses the three key value problems for knowledge, as well as the epistemic version of the Swamping Problem, than Pritchard's promising anti-luck virtue epistemology account.

There are some significant limitations to my account here, however. There is a wide range of problems in epistemology that my account has yet to fully address. For instance, Pritchard's anti-luck virtue epistemology explains why we don't think of cases like Gettier's cases and barn cases as representative or typical cases of knowledge. Pritchard claims that this is because there are two primary intuitions about knowledge: an ability intuition and an anti-luck intuition. It is possible that the Kantian account of knowledge I've offered could include some sort of condition that would better address these intuitions as well. However, it would need to be shown what such a condition might look like and why such a condition as a component of knowledge is not merely ad hoc. But the account I've offered here does deal with some of the luck problems, and is open to other necessary or sometimes present components of knowledge.

One other possible point of contention might regard the doxastic voluntarism of my account. Critics of doxastic voluntarism will question to what extent it makes sense to think that we can decide what to believe and effectively will to change our beliefs. I presented some examples above concerning at least some situations in which it seems possible to exercise direct voluntary control over our beliefs. And my account

specifically shows that epistemic duty discourages epistemic agents from exercising voluntary control over their beliefs for the wrong kind of reasons (inclinations, etc.). Even if we cannot have direct control over what or whom we come to believe on what grounds, we can exercise gradual control to some extent to acquire the habits of mind that epistemic duty consists in. In this way, epistemic agents can be praiseworthy (credit-worthy) or blameworthy in regards to their voluntary and willful compliance with epistemic duty. My account of knowledge offers a novel understanding of doxastic voluntarism that I would like to further explore.

My Kantian analysis of knowledge has some pretty interesting implications, as well. For instance, my account might mean that we need to reinterpret what gives rational beings like humans special dignity and thus special moral standing. Kant thinks that humans have a special dignity in their rationality, and thus some kind of special value. Neo-Kantians interpret the special dignity of humanity to reside in the capacity for moral rational choice and rational agency. My account, however, recognizes the unity of the structure of practical and theoretical reason, and thus indicates that we should reexamine this claim. One goal for further research might include fleshing out an argument that we should read the special dignity of humans as rationality, full stop, which includes not only rational choice, but also the kind of rational judgment epistemology deals with. This argument might conclude that beings who can think things for (the right kind of) reasons have dignity, and thus moral value.

Lastly, since Kant's moral theory includes an understanding of virtue as "the moral strength of a human being's will in fulfilling his duty" (6:405), a Kantian analysis of knowledge such as mine might be open to some appealing parts of various virtue

ethics accounts. Kant shows an interest in both moral virtues as strength of fulfilling one's moral duty and non-moral virtues that are at very least of instrumental and/or prudential importance in overcoming obstacles such as contrary interests or desires (see, for example, *Anthropology from a Pragmatic Point of View*). This account gives a good explanation as to why we think of open-mindedness and independent thought as important to the pursuit of knowledge (see the discussion on the maxims of 'common human understanding' above).

My analysis of knowledge offers an exciting way of understanding the goals and duties of ethics and epistemology as overall integrated and unified. Further explanation is needed of why, if these pursuits can be unified, why our epistemic goals might sometimes seem to contradict our moral goals. There is clearly more work to be done, but the novel approach offered here might serve to reinvigorate the search for answers that meet the challenges to our understanding of knowledge and its value.

BIBLIOGRAPHY

Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis.* Oxford University Press, Oxford, 1991.

Kant, Immanuel. *Grounding for the Metaphysics of Morals.* Translated by James W. Ellington, 3rd edition. Indianapolis, IN: Hackett Publishing Company, 1785/1993.

⸻. *Critique of Pure Reason.* Translated/edited by P. Guyer and A. Wood. Cambridge: Cambridge University Press, 1997.

⸻. *Lectures in Logic.* Translated/edited by J. Michael Young. New York: Cambridge University Press, 1992: 563-564.

⸻. *Practical Philosophy.* Translated/edited by Mary J. Gregor. New York: Cambridge University Press, 1996. Includes *Critique of Practical Reason* and "What is Enlightenment?"

⸻. *Religion and Rational Theology.* Translated/edited by A. Wood and G. di Giovanni. Cambridge: Cambridge University Press, 1996. Includes "What Does it Mean to Orient Oneself in Thinking?"

⸻. *Anthropology, History, and Education.* Translated/edited by G. Zoller, R. Louden, M. Gregor and P. Guyer. Cambridge: Cambridge University Press, 2007. Includes the *Anthropology.*

O'Neill, Onora. "Vindicating Reason," in *The Cambridge Companion to Kant.* Edited by P. Guyer. Cambridge: Cambridge University Press, 1992: 280-308.

Pritchard, Duncan; Millar, Alan; and Haddock, Adrian. *The Nature and Value of Knowledge: Three Investigations.* Oxford: Oxford University Press, 2010.