

Epistemic Supererogation and Its Implications

Trevor Hedberg

This version of the article is the penultimate version that was accepted for publication in *Synthese*. Those wishing to cite the paper must consult the published version. The final publication is available at Springer via <http://dx.doi.org/10.1007/s11229-014-0483-5>.

1 Introduction

Supererogatory acts, those which are praiseworthy but not obligatory, have become a significant topic in contemporary moral philosophy, primarily because morally supererogatory acts (MSAs) have proven difficult to reconcile with other important aspects of normative ethics. However, despite the similarities between ethics and epistemology, epistemic supererogation has received very little attention. Alvin Plantinga briefly mentions the concept on a few occasions (Plantinga 1986, p. 7; 1988, p. 10; 1991, p. 296), and Paul Tidman provides a more thorough treatment (Tidman 1996, pp. 269-270, 274). But in each case, the examination of epistemic supererogation is limited to (at most) a few paragraphs, and no elaborate treatment of the subject exists in the literature. As a result, the implications of epistemically supererogatory acts (ESAs) remain largely unexplored.

One might think that the lack of literature on the topic is explained by the fact that actions cannot be epistemically evaluated *at all*. A proponent of this view would claim that epistemic evaluations apply only to cognitive states. On such an account, if anything is epistemically supererogatory, it must be certain cognitive states because such an epistemic evaluation could

not properly apply to any actions. The proponent may also claim, as Richard Feldman does, that epistemic duties apply only to how we use our mental faculties and not to any actions we perform (Feldman 2002). Supererogation is usually thought to identify actions that go beyond the call of duty. If this colloquial description is remotely accurate, then it is hard to see how we could perform epistemically supererogatory acts unless there were some epistemic duties pertaining to actions.

I cannot argue for it at length here, but an assumption of this paper is that there are some actions that can be epistemically evaluated and that there are some epistemic duties that pertain to actions.¹ Following Hilary Kornblith (1983), I believe that the notion of epistemic responsibility provides a basis for epistemically evaluating one's actions. We can certainly speak of epistemic justification with respect to how well a person reasons (e.g., whether or not they adopt the attitude consistent with the evidence available to them), but this is not the *only* kind of epistemic justification that is worth discussing. As Kornblith puts the point, "Sometimes when we ask whether an agent's belief is *justified* what we mean to ask is whether that belief is the process of *epistemically responsible action*, i.e. the product of action that an epistemically responsible agent might have taken" (Kornblith 1983, p. 34). Kornblith defines an epistemically responsible agent as one who (1) desires to have true beliefs, and therefore desires her beliefs to be formed through processes conducive to forming true beliefs, and (2) acts in accordance with these desires. Examples of epistemically responsible action are not hard to produce. Among others, they include properly gathering evidence, paying due attention to reasonable criticisms of one's views, and guarding against defects in one's reasoning (e.g., self-deception, perceptual

¹ Prior treatments of epistemic supererogation also assume that epistemic evaluations can apply to actions. Plantinga speaks of "supererogatory effort" to live a life of epistemic excellence (Plantinga 1986, p. 7; 1988, p. 10). Such a pursuit is an activity that one undertakes. Tidman argues that certain kinds of critical reflection are epistemically supererogatory, and critical reflection is a mental action that we can voluntarily perform.

illusions). The problem with viewing epistemic justification as concerning only cognitive states is that such a construal is too narrow; some of our actions also play a role in whether our beliefs are justified. I contend that we can be epistemically praised and epistemically blamed for these actions. Furthermore, since epistemically responsible action is tied to justified belief, I believe there can be epistemic duties based on acting in epistemically responsible ways.

In this paper, I aim to further the discussion of how we can epistemically evaluate actions by arguing for the existence of ESAs and considering the potential implications of their existence.² First, I offer a brief account of MSAs and how their existence generates a strong intuition that a similar phenomenon should exist in epistemology. Afterward, I examine five cases of epistemic activities that appear to be epistemically supererogatory. I argue that deeming these acts epistemically supererogatory is the best explanation for our considered verdicts about the different cases: in each case, the individual appears to perform an activity that is epistemically praiseworthy but (for various reasons) cannot be epistemically required of them. Finally, I consider how the existence of ESAs might impact the contemporary study of epistemology, particularly with regard to our conceptions of epistemic duties.

2 Moral Supererogation

In the history of moral philosophy, many once thought that there were only three types of moral actions: required actions, forbidden actions, and merely permissible actions (i.e., those that are not required, forbidden, or supererogatory). Generally, required acts are morally good to perform, and forbidden acts are morally bad to perform.³ Admittedly, not all ethical traditions

² Although it will not be my focus, I will return to the possibility of there being epistemically supererogatory cognitive states in the concluding section.

³ There can be cases where a morally good action is forbidden because better alternatives are available, and there can also be cases where a morally required action is morally bad because all the alternatives available are awful.

embraced the tripartite schema,⁴ but explicit discussion of supererogation in non-religious ethical theory nevertheless appears to be a relatively recent phenomenon, emerging when J.O. Urmson (1958) presented MSAs as a challenge to the tripartite classification scheme. Urmson posited actions performed by saints and heroes as paradigm examples of those that are morally praiseworthy but not morally required, and his arguments initiated a widespread debate in moral philosophy.⁵ Some criticized Urmson's arguments and resisted acknowledging the existence of MSAs,⁶ but this dissent has gone against the prevailing trend. A mere 33 years after Urmson's "Saints and Heroes" was published, Susan Hale described the new fourfold categorization of moral acts (required, forbidden, merely permissible, and supererogatory) as "near dogma" (Hale 1991, p. 273).

Perhaps the reason for this trend is that certain paradigm examples of MSAs seem impossible to properly describe under the tripartite classification system. Imagine an act of extreme beneficence: late in life and without any living relatives, a person voluntarily gives all her life savings to charity. It would be ludicrous to claim that this action was morally forbidden: we commend this action rather than condemning it. Simultaneously, it seems absurd to claim that she was morally required to make such a sacrifice: she could have permissibly spent some of those earnings on herself (e.g., purchased books, gone to the theater). After all, we do not ordinarily condemn those who keep some of their lifetime earnings to themselves, even late in their lives. We seem forced to conclude that the action is morally permissible, but unlike standard (merely) permissible actions, this one is morally praiseworthy. Hence, a fourth category of moral actions seems necessary to fill this conceptual gap. Supererogatory actions are not

⁴ The Roman Catholic ethical tradition is a notable example.

⁵ For useful surveys of the literature written in response to Urmson's article, see Jackson (1986) and Stanlick (1999).

⁶ The definitive critic of Urmson in this regard is Elizabeth Pybus, who defends the traditional trichotomy of moral actions. See Pybus (1982).

necessarily limited to acts of extreme beneficence, but these kinds of cases seem to be the hardest to dismiss. Beyond charitable giving, acts of heroism – such as a bystander’s voluntary attempt to save others trapped in a burning building – are also strong candidates for supererogatory behavior.

Given the ongoing debate in moral philosophy, I doubt everyone will find appeals to paradigm examples of MSAs convincing, but elaborate discussion of MSAs and their conceptual coherence is not my primary task here. For the duration of the paper, I assume that MSAs (in some form) exist. This concession, though noteworthy, is not a substantive weakness in my argument for two reasons. First, the trend in contemporary moral philosophy is not to defend the existence of MSAs but to make them compatible with our general moral principles or specific normative ethical theories.⁷ Second, my argument for the existence of ESAs does not rely *entirely* on the existence of MSAs for its support. The presence of MSAs does lend some credence to the existence of ESAs since ethics and epistemology are isomorphic frameworks in some important respects (a thought I develop further in the next section), but the existence of MSAs is not strictly required for my argument to succeed. Thus, even those skeptical of the existence of MSAs should consider the possibility of ESAs.

Before proceeding, one further bit of clarification is in order. While I have alluded to the standard features of MSAs in the prior analysis, it will be helpful to provide a precise definition, especially since my definition of an ESA will parallel my definition of an MSA. Following Gregory Mellema’s account of moral supererogation (Mellema 1991), an act is an MSA iff (1)

⁷ To give a few general examples, see Attfield (1979), Jackson (1986, 1988), and Postow (2005). Attempts to accommodate MSAs into certain moral theories have been undertaken by both consequentialists (Sikora 1979, Vessel 2010) and Kantians (McCarty 1989). Certain debates also presuppose that MSAs exist. For instance, when one considers whether or not a promise to perform an MSA is obligatory (e.g., Kwall 2005), one must assume that there are some genuine MSAs. Otherwise, the worrisome paradox (about being obligated to perform a supererogatory action) could not arise, since there would be no MSAs to promise to perform.

performing the act fulfills no moral duty, (2) performing the act is morally praiseworthy, and (3) omitting the act is not morally blameworthy.⁸ Now we are in position to consider what relevance MSAs have to contemporary epistemology.

3 Ethics and Epistemology as Isomorphic Frameworks

Epistemologists often draw upon contemporary work in ethical theory to develop their views on subject matter in epistemology. Epistemic duties are sometimes described as analogous to moral duties, and work in virtue epistemology frequently derives from work in virtue ethics. In some cases, more specific parallels are drawn. Reliabilism, for example, is sometimes considered the epistemic equivalent of rule utilitarianism (Firth 1981, Zagzebski 1996).

The interactivity of ethics and epistemology seems natural for several reasons. First, they are each normative frameworks: ethics concerns what we morally ought to do, and epistemology concerns what we epistemically ought to believe. Moreover, as discussed in previous sections, part of determining what we epistemically ought to believe involves acting in epistemically responsible ways. As a result, we can sensibly discuss actions that we epistemically ought to do; both ethical duties and epistemological duties specify how we ought to act. Second, these disciplines overlap in the field known as the ethics of belief, the sub-discipline which focuses on the moral considerations pertaining to how we form beliefs and the relationship between moral normativity and epistemic normativity. Some even contend that epistemic normativity is ultimately grounded in moral normativity⁹ (implying dependence of the former upon the latter), although many think that these frameworks interact in a less overt way. For our purposes here,

⁸ I assume in this paper that one is never morally permitted to do something morally blameworthy and that one is never morally forbidden from doing something morally praiseworthy.

⁹ One of the classic presentations of this idea comes from Clifford (1879). For a more contemporary account of this view, see Grimm (2009).

we can avoid dwelling on the varying opinions of epistemologists in this regard. We only need to note that epistemology and ethics are similar disciplines in some important ways and that they are not typically thought to be wholly independent of one another.

At this juncture, we can see why the existence of epistemically supererogatory acts (ESAs) appears intuitively plausible. Both are normative frameworks, and there is some overlap between them. Thus, to facilitate coherent accounts of these disciplines, we have some incentive to conceive of them as isomorphic entities.¹⁰ The influence of contemporary work in moral philosophy on epistemology furthers the case for conceiving of the frameworks in this manner. Hence, if we conceive of ethics and epistemology as isomorphic, and if certain types of acts can be identified as morally supererogatory, then we should expect to find acts which can be identified as epistemically supererogatory. While appealing to the parallelism of ethics and epistemology does not provide us an argument that ESAs exist, it does provide a good reason to investigate the matter. Perhaps we will discover that ethics and epistemology are dissimilar when it comes to supererogatory acts, but we have a good *prima facie* reason to think this will not be the case.

4 Unpacking Terminology and Clarifying Concepts

To keep things simple and stay consistent with the isomorphism of ethics and epistemology, I define an ESA in a parallel fashion to an MSA. Specifically, an act can be considered an ESA iff (1) performing the act fulfills no epistemic duty, (2) performing the act is epistemically praiseworthy, and (3) omitting the act is not epistemically blameworthy. I follow Richard Feldman in thinking that epistemic duties “are the duties we must carry out in order to

¹⁰ There may, of course, be some minor dissimilarities between the two. The isomorphism between ethics and epistemology is suggested as a general rule rather than an absolute one.

have justified beliefs” (Feldman 2002, p. 380).¹¹ In the context of this definition, the relevant epistemic duties are those that apply to actions: they are those actions that we must perform to live in epistemically responsible ways. In section 1, I gestured at some examples of epistemically responsible actions, but an example will help illustrate how these actions are linked to my definition of epistemic duties.

Imagine that Jones, a young physicist, is presenting a paper and desperately eager to hear the praise of his colleagues regarding his research, but a senior colleague presents a cogent, devastating criticism. Jones, unable to tolerate criticism, pays no attention to the objection that is raised, and so his beliefs do not change. By refusing to even listen to the criticism, Jones violates an epistemic duty because he acts in an epistemically irresponsible manner.¹² His action is inconsistent with a genuine desire to have true beliefs, and this deliberate avoidance of acknowledging criticism will (most likely) lead Jones to form false beliefs rather than true ones. Thus, Jones’ beliefs about the merits of his research are unjustified because they are maintained in part by a behavior (i.e., willfully ignoring criticism) that is epistemically irresponsible. In this case, we can identify a plausible candidate for an epistemic duty: one should be attentive to reasonable criticisms of the beliefs that one holds. While there are many others that could be posited (including minor alterations of the one just proposed), unpacking a full list of these kinds of epistemic duties is beyond the space available here. However, it is worth noting that these duties do not require epistemic perfection. A duty to be attentive to criticism, for instance, does not require that one be unfailingly alert to all possible criticisms that could be raised against her

¹¹ As mentioned earlier, Feldman thinks that there are no epistemic duties pertaining to actions. My disagreement with Feldman on this point arises from his claim that the *only* requirement for having a justified belief is that one adopts the attitude about a given proposition that is justified by the evidence presently available regarding that proposition. I believe that whether one’s beliefs are justified is a more complicated matter than this. To give one illustration, beliefs formed by deliberately avoiding the acquisition of counterevidence are surely sometimes unjustified, even though the cognitive attitude that one adopts regarding the relevant propositions may perfectly mirror the evidence that one has access to at the time.

¹² I borrow this example from Kornblith (1983, p. 36).

beliefs. Such a standard might be appropriate for God or other ideal epistemic agents, but it would be ludicrously unrealistic to set such lofty requirements for ordinary human beings. (If we did set our standards that high, no human beings would be epistemically responsible agents.) A duty to be attentive to criticism is better understood as a duty to acknowledge and seriously consider criticisms that are substantive and articulated clearly. This epistemic duty – like others that we could examine – specifies a threshold that must be met for one’s conduct to be epistemically responsible. Beliefs formed or maintained through epistemically irresponsible conduct are unjustified.

Moving on to conditions (2) and (3) of ESAs, in the context of this definition, epistemic praise and blame are (like epistemic duties) tied to actions that agents perform. An agent is epistemically praiseworthy for performing an action iff that action violates no epistemic duties and can be reasonably expected to bring about more epistemic benefits than epistemic harms. These epistemic benefits could include (among other things) the acquisition of true beliefs, the elimination of false beliefs, the improvement of a particular cognitive attitude (e.g., turning a mere true belief into knowledge), or the cultivation of an epistemic virtue (e.g., open-mindedness).¹³ An agent is epistemically blameworthy for performing an action iff that action violates an epistemic duty or can be reasonably expected to bring about more epistemic harms than epistemic benefits.¹⁴ Epistemic harms could include the acquisition of false beliefs, the

¹³ I cannot offer a comprehensive list of epistemic virtues and vices, but a robust list of that sort will not be necessary for the case analyses in the next section.

¹⁴ In defining epistemic praise and blame in this manner, I leave open the possibility that there are cases where one is epistemically blameworthy for performing an action even though the action is not epistemically irresponsible (and so violates no epistemic duty). In this manner, my definition of ESAs allows for the possibility that failing to perform a particular action could be epistemically blameworthy even if performing the action is not required by any epistemic duty, which explains why conditions (1) and (3) are not redundant. This is parallel to how Mellema’s account of MSAs allows for cases where an act’s omission could be blameworthy even if there is no moral duty that the act be performed.

elimination of true beliefs, the regression of a cognitive attitude (e.g., a justified true belief becoming a mere true belief), or the cultivation of an epistemic vice (e.g., closed-mindedness).

Urmson initiated thorough examination of moral supererogation by presenting examples of acts which he thought could only be classified as morally supererogatory, and I will echo his strategy. In what follows, I argue for the existence of ESAs by presenting five cases in which I believe the agent's action can only be properly classified as epistemically supererogatory. The most plausible and paradigmatic instances of MSAs are those in which an agent can provide altruistic gains at nontrivial personal costs. In these cases, respect for the agent's autonomy renders either option permissible; providing the altruistic gains is morally supererogatory. In a sense, all of my cases will follow a parallel structure: agents will have the ability to generate epistemic benefits through their conduct but only at some personal cost (e.g., time, energy). In some of the cases (particularly the second and fourth cases), the personal costs are significant and concern for agent autonomy seems to be the key element in explaining why the agent's action is epistemically supererogatory, but in others, it is not clear the extent to which the agent's behavior is a particularly costly sacrifice. This fact shouldn't worry us too much, however, because we need not think at the outset that the analogy between moral supererogation and epistemic supererogation will hold strictly in every respect.

Another asymmetry between MSAs and ESAs may lie in the scope of merely permissible actions. One may think that overall balance of epistemic benefits and epistemic harms produced by an action will only rarely be equivalent, and so only a small number of actions will, according to my definitions, be *merely* epistemically permissible – allowable and neither praiseworthy nor blameworthy. In the moral domain, merely permissible actions seem fairly common, and thus, one may find their rarity in the epistemic domain a bit striking. However, I am not certain that

merely permissible actions are in fact so rare in the epistemic realm. There may be cases where the epistemic benefits and harms generated by an action are incommensurable with one another: in such cases, there would simply be no fact of the matter about whether the epistemic benefits were greater than, less than, or equivalent to the epistemic harms.¹⁵ Provided that this kind of action neither fulfills nor violates an epistemic duty, it might be described as merely permissible, since it would not be epistemically praiseworthy or blameworthy according to my definitions. Nevertheless, it must be acknowledged that the precise scope of (merely) epistemically permissible actions must remain somewhat unspecified throughout this paper. The relative scarcity of these actions (if indeed these kinds of actions prove to be scarce) and the concerns generated by their scarcity are issues that would need to be addressed in future work on ESAs, work in which the relationship between ESAs and (merely) epistemically permissible actions is the primary focus. My task here is only to establish that ESAs exist and to consider a few of the potential implications of their existence. With these aims in mind, we can now turn to some examples of ESAs.

5 Paradigm Examples of ESAs

For our first case (C1), let's consider Tim and his debit card. Tim has memorized the pin number associated with his card and has never misremembered it. Nevertheless, he occasionally checks his pin number to be sure he has it correct. It seems safe to say that Tim has no epistemic

¹⁵ For an overview of the philosophical issues surrounding incommensurability, see Hsieh (2007). The application of this concept to the issue at hand is that there can be cases where an action generates epistemic benefits and harms that are simply too different to be compared to one another. Some cases certainly admit of sensible comparisons between epistemic harms and benefits. When an action produces no effects other than two true beliefs, it clearly produces more epistemic benefits than epistemic harms. But this kind of comparison can easily become more complicated. For instance, does the epistemic harm of forming a false belief outweigh the epistemic benefits of turning two true beliefs into justified true beliefs? Does the epistemic benefit of coming to understand a subject outweigh the epistemic harms of forming several false beliefs about a different subject? These kinds of questions do not admit of easy answers, and in at least some cases where actions produce effects like these, the epistemic benefits and harms may prove incommensurable with one another.

duty to check his pin number: he has sufficient evidence for believing his memorized number is correct, which justifies his belief. Moreover, if Tim were not to engage in this activity, surely we would not claim he was epistemically blameworthy for acting in this way: there is nothing epistemically irresponsible about refraining from double-checking beliefs that are supported by sufficient evidence, unless perhaps counterevidence is present (which is not the case here), and such behavior does not ordinarily lead to epistemic harms unless perhaps one is especially forgetful.¹⁶ Furthermore, Tim seems epistemically praiseworthy for acting in this manner. Tim's act of cognitive double-checking increases the certainty that his memorized number is correct and decreases his chances of misremembering it. Admittedly, Tim's act of checking his pin number could lose its praiseworthy status if he did it needlessly or obsessively (such that it seemed like an irrational compulsion), but let's suppose that this act is done only occasionally and serves the epistemically useful purpose of helping Tim retain a true belief. Under these circumstances, conditions (1)-(3) are met, and Tim's act appears to be an ESA. Considering that these types of actions (i.e., occasional double-checking of epistemically justified and uncontroversial beliefs) seem fairly common, ESAs may not be a particularly rare phenomenon.

But this analysis might seem too quick. We should consider a potential objection: one might try to undercut C1 by challenging whether (2) is met. The objector could contend that Tim is engaging in exceptionally *prudent* cognitive behavior and that this does not seem equivalent to epistemically supererogatory behavior. We can imagine cases in which exceptionally prudent cognitive behavior does not seem epistemically praiseworthy. Imagine that Tim is placed in an

¹⁶ My definitions of epistemic blameworthiness and epistemic praiseworthiness both require that the actions in question can be "reasonably expected" to bring about certain effects. This clause is meant to account for bizarre circumstances in which the epistemic consequences resulting from certain actions are much different than what could ordinarily be expected. Throughout these case analyses, when discussing the epistemic benefits and harms caused by the actions under scrutiny, I will generally assume that the consequences under discussion are consistent with what we could reasonably expect to result from these actions. Hence, explicit examination of reasonable expectations about the actions' consequences is generally absent.

extremely high-stakes case in which he engages in obsessive double-checking over an extremely simple and obvious proposition. His action may be prudent, but it is not epistemically praiseworthy. Hence, reflection on C1 does not justify belief in ESAs.

While the objector's concern is understandable, we can use one of her assumptions to construct an alternative case which better supports the existence of ESAs. The objector implicitly differentiates between two types of praise: prudential and epistemic. In C1, we are purportedly praising Tim's action on prudential grounds rather than epistemic ones. Hence, if these two types of praise can diverge, then we can construct an alternative case which inverts the nature of praiseworthiness. Suppose Tim engages in periodic double-checking about a particular mathematical formula. This piece of information has no obvious usefulness to Tim beyond its status as a true belief. In this case, it seems we can actually criticize Tim on prudential grounds: surely he has something better to do with his time than double-check the truth of a trivial belief. On epistemic grounds, however, we can praise Tim for engaging in (non-obligatory) behavior which helps him retain a true belief. We may find that the *amount* of praise for this type of belief retention is much smaller than in the original C1, but this discovery should not be surprising. In C1, we find Tim praiseworthy in *two* ways at the same time: he is praiseworthy on both epistemic and prudential grounds. In this modified case, Tim is only epistemically praiseworthy and may actually be prudentially blameworthy (depending on what else he could be doing). Nevertheless, being even *a little* epistemically praiseworthy is all that is required for being an ESA. There is no reason to think that all ESAs must be epistemically praiseworthy to a great extent since epistemic praiseworthiness can presumably come in degrees.

Our second case (C2) is modeled after an argument made by Paul Tidman regarding critical reflection (Tidman 1996, pp. 169-170). Imagine that John decides to spend a portion of

his evening critically reflecting on many of the beliefs he has formed subconsciously throughout the day. These include beliefs about what his wife had for breakfast, why a colleague was leaving the office early, where a friend ate lunch that day, and so on. In reflecting on these beliefs, John will likely be able to identify some beliefs which are false or based on poor evidence. After discarding these beliefs or withholding judgment about them until he can gather further evidence, he will be in an epistemically superior cognitive state. Consequently, this activity yields epistemic benefits without yielding any notable epistemic harms, rendering John epistemically praiseworthy for performing the action. However, it does not seem to be an epistemic duty.

An epistemic duty requiring critical reflection on all of our day-to-day beliefs would be too demanding to fulfill, and even a less stringent duty could be failed through no fault of one's own. Tidman offers the example of believing incorrectly that he saw a friend earlier in the day (Tidman 1996, p. 270). Were he to reflect on the belief, evidence (e.g., physical appearance of the person, the knowledge his friend was out of town that day) would lead him to realize that the person he saw was not actually his friend. However, he is distracted by a nearby car accident before he can engage in this reflection, and the thought never occurs to him again. It seems absurd to suppose that he failed an epistemic duty in not reflecting on this belief. Therefore, Tidman concludes that we do not have a duty to reflect on mundane, everyday beliefs in a rigorous manner. Insisting that one must always critically reflect on such beliefs is "neither realistic nor reasonable" (Tidman 1996, p. 270). From this observation, we can also infer that epistemically responsible agents would not always critically reflect on everyday beliefs like John does. Additionally, unless John deliberately avoids ever critically reflecting on beliefs (cultivating an epistemic vice in the process), a failure to critically reflect on these particular

beliefs does not seem to bring about more epistemic harms than benefits: not performing this action would neither help nor harm his cognitive state. Thus, John would not be epistemically blameworthy for failing to perform this action. Hence, John's act of critical reflection can be described as an ESA because it meets (1)-(3).

After considering C2, one might object that (1) is not met. While we may not have a duty to critically reflect on *all* our subconsciously acquired mundane beliefs, we do not need to endorse such a strong claim to see John as fulfilling an epistemic duty. Suppose we have a relatively weak (*prima facie*) duty to frequently set aside brief periods of time to consider the epistemic credentials of a sampling of recently formed beliefs. This activity might help ensure that we do not become too cavalier about forming beliefs and curtail our believing propositions on insufficient evidence. On this (much more plausible) account of the duty to critically reflect, John's action fulfills an epistemic duty and is not an ESA.

C2 can easily be modified to accommodate this objection. Suppose that this modified duty of critical reflection requires (on average) ten minutes of daily critical reflection. Now suppose John engages in the activity of critical reflection for an average of fifteen minutes per day.¹⁷ Because John's actions go beyond what duty requires and also fulfill (2) and (3), they are epistemically supererogatory.

While both C1 and C2 pertain to types of critical reflection on beliefs, we should note that they concern different types of ESAs. C1 pertains to double-checking the truth of epistemically justified beliefs which are uncontroversial to increase the certainty of their truth and reduce the chances of forgetting them. In contrast, C2 pertains to critically reflecting upon beliefs which are acquired subconsciously and would ordinarily go unexamined.

¹⁷The amount of time chosen here is arbitrary and inconsequential. However long the duty of critical reflection requires John to reflect on his beliefs, we can posit that John performs the activity for longer than the duty requires.

Now let's turn to a third case (C3): Anna is researching the debate in metaphysics between realists, nominalists, and idealists regarding the problem of universals.¹⁸ Anna reads 10 articles and has sufficient evidence to conclude that realism is the correct view. However, she notices another pertinent article in her pile of research materials that she has not read. She elects to read it before evaluating her evidence and forming a belief. Let us suppose that Anna has gathered the evidence and examined it in epistemically responsible ways. (She has not, for instance, gathered and read 10 articles from the same author.) Under such circumstances, it seems that after reading the 10 articles, Anna has already fulfilled what her epistemic duties could require. But was Anna's reading this additional article an ESA?

Based on the remarks above, (1) is met. Additionally, (3) appears to be met: Anna would not be violating an epistemic duty or bringing about any overall epistemic harm by electing not to read this last article. However, it is less clear whether (2) is met. In fact, some might argue that Anna is epistemically *blameworthy* for reading this additional article before making a judgment. One can be *too* skeptical, and since Anna has sufficient evidence for forming a belief, she ought not to reserve judgment in this case. Her skepticism here might represent an epistemic defect rather than epistemic excellence. However, while I agree that radical skepticism should not be construed as epistemically praiseworthy, I think this interpretation of Anna's action is uncharitable.

If reading this additional article was wholly frivolous, then perhaps it is not epistemically praiseworthy and could be deemed a waste of time, but it would be unreasonable to suppose that Anna gains nothing (epistemically speaking) from reading the article. While it is doubtful (based

¹⁸ While many topics could be used in C3, I have chosen this one in an attempt to avoid the encroachment of moral praise. If the issue presented were global climate change, for example, the praiseworthiness of Anna's actions could be derived from the moral relevance of the topic she was researching. My hope is that this debate in metaphysics is regarded as morally neutral.

on the previous articles she has read) that reading the eleventh article will radically alter her view, it will almost surely provide her some additional knowledge about the topic. Furthermore, even if it only reiterated points made in previous articles, it would further instantiate those points and give them greater support. Moreover, we need not interpret her reservation in withholding judgment as a cognitive failure. Perhaps she withholds judgment so that the attitude that would result cannot affect her reading of this last article.¹⁹ Of course, at some point, this practice would be foolish and not worthy of praise: if Anna decided to read 50 additional articles before endorsing a view on the topic, then her reservations would be difficult to justify, and her activities would probably demonstrate an epistemic deficiency.²⁰ Anna's action in C3 is not nearly this problematic, however, so we can safely deem Anna epistemically praiseworthy for performing this act. Hence, (2) is met, and Anna's act is an ESA.

In our fourth case (C4), let's consider an ESA that increases one's stock of true beliefs. Suppose Wendy decides to devote an hour each night to reading an encyclopedia instead of watching reruns of her favorite television series. In performing this activity, Wendy increases the

¹⁹ Anna may, for example, worry that her attitude about the topic would affect how she perceived the evidence presented in the article. When confronted with information that appears inconclusive, people will almost always interpret the information in a way that supports beliefs they already hold. This psychological phenomenon is often called "Confirmation Bias." One of the classic studies that discusses this phenomenon is Lord, Ross, and Lepper (1979). A useful overview of some relevant literature on the topic can be found in Gilovich (1991). Notably, Thomas Kelly has recently written about this phenomenon in an epistemic context; see Kelly (2008).

²⁰ One may worry about a Sorites paradox here. Anna's decision to read 50 additional articles on the topic may indeed demonstrate unwarranted skepticism, but will she still manifest this unwarranted skepticism if she decides to read an additional article 50 times in succession? If she examines an additional article and decides to read it before forming a judgment, and then reiterates this process 49 more times, one may think that her action is epistemically praiseworthy each time, which means that her reading of 50 additional articles would be epistemically praiseworthy if carried out in this manner. But that outcome clearly looks incorrect: surely Anna cannot be praiseworthy for reading 50 additional articles and being so epistemically overcautious. Fortunately, if we keep the details the same as specified in C3, I think this worry can be allayed. Remember that the initial 10 articles provided sufficient evidence for Anna to conclude that realism was the correct view. If she decides to suspend judgment and read an extra article, she is epistemically praiseworthy for performing this particular action. She might also be praiseworthy for reading a second article – perhaps even a third. But with each subsequent article that she reads, Anna gets closer and closer to an unhealthy skepticism that is epistemically blameworthy: at a certain point, the weight of the evidence will become so great that it become epistemically irresponsible of her to continue suspending judgment. Of course, it may not be possible to specify precisely where Anna's decision to read an additional article becomes blameworthy (rather than praiseworthy), but this only shows that the progression from reasonable further inquiry to unwarranted skepticism is a gradual one.

number of significant true beliefs she holds. Even if Wendy were to form the same total number of beliefs if she had watched this show (e.g., beliefs about what she was doing at the time), these kinds of beliefs do not seem as epistemically valuable as those that would be acquired from encyclopedia reading or similar scholastic activity. No epistemic duty to engage in this activity seems present, unless there is some requirement that we be nearly always striving to increase our stock of significant true beliefs. Such a requirement would surely be too demanding: one would think that epistemic duties, like moral duties, can only reasonably require a certain degree of self-sacrifice.²¹ Wendy's action also can be reasonably expected to bring about significant epistemic benefits without generating any epistemic harms, and so she seems epistemically praiseworthy for performing this action. A failure to perform this action also would not be blameworthy: omitting this action would not bring about more epistemic harms than epistemic benefits, and our epistemic duties do not unceasingly on the lookout for ways in which we can increase our stocks of true beliefs. Since (1)-(3) are satisfied, Wendy's action is an ESA. Similar to C1, cases resembling C4 abound in the real world. We frequently have the opportunity to increase our knowledge when we have no epistemic duty to do so. After all, our time is limited, and we cannot be required to perform knowledge-gathering activities every moment we are not preoccupied. In circumstances where we sacrifice our leisure for the pursuit of knowledge when an epistemic duty does not require us to do so, we are often performing ESAs.

For our fifth and final case (C5), suppose Adam is investigating some issues in quantum physics. He discovers that experts in the field have, with few exceptions, endorsed a particular unified field theory—call it Theory X. Adam has a very limited understanding of quantum physics and would be justified in accepting the consensus of experts on the matter. But suppose

²¹ I allude here to the general thought that a moral theory that forces an agent to make extraordinary personal sacrifices is implausible. This objection is commonly lodged against utilitarianism.

Adam wants firmer grounding for this belief and proceeds to investigate it on his own, reading some articles written by the experts and educating himself about the subject matter so that he can draw his own conclusion about Theory X. This activity might alter Adam's epistemic state from one of justified belief or justified true belief to one of genuine knowledge (or even understanding). However, it is surely too bold to say this type of investigation is required of all controversial topics outside our areas of expertise: the demandingness of such a claim is absurd. It must sometimes be epistemically permissible to appeal to the consensus of experts, especially with regard to topics beyond our understanding, and appealing to experts seems perfectly consistent with what an epistemically responsible agent would do. So Adam's action is not required by any epistemic duty, and Adam is praiseworthy for acting in this way because his action brings about epistemic benefits without any corresponding epistemic harms. A failure to act in this manner also would not be epistemically blameworthy because appealing to the consensus of experts (rather than investigating the issue independently) cannot be reasonably expected to result in a greater balance of epistemic harms than benefits. Thus, this act meets (1)-(3) and should be deemed an ESA.

While I doubt that most readers will accept all five cases as examples of ESAs, one does not have to accept my analysis of *all* the cases for the basic argument to succeed. If my analysis of even one case is convincing, then ESAs exist. I also suspect that there are other types of ESAs which would not fit the paradigms expressed in C1, C2, C3, C4, or C5.²² Because the concept of epistemic supererogation has been largely unexplored, it remains unclear how many types of acts could be properly described as epistemically supererogatory.

²² We may, for instance, be able to construct a case in which one is plausibly under no duty to gather counterevidence but is epistemically praiseworthy for doing so.

6 The Implications of ESAs

A rigorous analysis of the implications of ESAs is beyond the scope of this paper, but it is certainly worth considering how the existence of ESAs may affect our understanding of epistemology. Primarily, epistemic supererogation provides a new criterion for evaluating principles pertaining to the ethics of belief, especially alleged epistemic duties. Specifically, we should reject any principle which *requires* individuals to perform acts which are better described as epistemically supererogatory. To illustrate how we can employ epistemic supererogation to evaluate such principles, let us examine two proposed epistemic duties.

First, consider Kihyeon Kim's claim that "deontological terms apply, first and foremost, to critical reflection on a belief: one *ought to* reflect on beliefs; one has an *epistemic duty* to reflect on them" (Kim 1994, p. 284, original emphasis). Tidman appeals to ESAs to argue against Kim's proposed duty of critical reflection on the grounds that this duty requires us to perform actions which are actually epistemically supererogatory. Kim's account does not impose any restrictions on what beliefs we are required to critically reflect upon: it seems we are required to critically reflect on *all* beliefs we hold. While Tidman acknowledges that critically reflecting on beliefs will help us eliminate false beliefs, he claims it is absurd to maintain that we are *required* to reflect on every ordinary belief we come to hold during each day. (Some details of his criticism are discussed in C2 in section 5: John's actions are designed to resemble those of a person who engages in critical reflection of the sort that Kim describes.) As Tidman summarizes, "Given the daunting number of new beliefs we acquire every day, such constant heroic efforts are beyond what a reasonable conception of epistemic duty can require" (Tidman 1996, p. 270). We may have duties to critically reflect on certain important or controversial beliefs we hold, but this duty cannot extend to *every* belief we hold. Thus, Kim's duty to

critically reflect should be rejected unless it can be modified in a way that makes it less demanding or suitably constrained by a supplementary principle.

Second, consider a proposed duty of evidence gathering. Richard Hall and Charles Johnson argue that we can have the following epistemic duty to seek more evidence (D): “For any proposition that is less than certain on one’s present evidence, one has an epistemic duty to seek more evidence about that proposition” (Hall and Johnson 1998, p. 133). Although Hall and Johnson assert (D) as a *prima facie* duty (rather than an absolute one), this still seems problematic. Suppose two undergraduate students approach me. One claims I have 15 pens in my desk drawer, but the other claims I have 16. Do I really have a *prima facie* duty to count the pens in my desk drawer as (D) suggests? If I were to sift through the drawer and count my pens, we might say that I deserve some epistemic praise for performing this action (since I’ll acquire a new true belief and incur no epistemic harms), but even if I had the time and resources to pursue this task (and it did not hinder my ability to perform some other more important task), could we really claim that I have failed an epistemic duty if I do not count the pens in my drawer? I cannot see how a failure to perform this act would violate any epistemic duty or how doing so would in any other way make me epistemically blameworthy. This act is better classified as an ESA. Since (D) *requires* us to perform ESAs, the principle must either be rejected or be modified to overcome this difficulty.

Although the existence of ESAs might have substantial repercussions regarding how we critique (at least some) epistemic duties, their implications may stretch deeper into epistemology than these illustrations indicate. ESAs may have some bearing what theory of epistemic justification proves most plausible. They may prove particularly relevant to some versions of reliabilism, an approach to epistemic justification that places an emphasis on how reliably one’s

method of obtaining a belief is truth-conducive. Reliabilism is a perfection-driven approach to epistemology, one that shares some similarities to consequentialism in the moral domain.

Consider how Linda Zagzebski describes these parallels between epistemology and ethics:

Since contemporary epistemology is belief-based, it is no surprise that the type of moral theory from which these theories borrow moral concepts is almost always an act-based theory, either deontological or consequentialist. So we generally find that epistemologists refine their inquiry into one of two types of questions: (1) does the belief violate any epistemic rules or any epistemic duties? Is it epistemically permissible, within one's epistemic rights? Theories of this sort take deontological moral theories as their normative model. (2) Was the belief formed by a reliable process for obtaining the truth? Theories of this sort are the forms of reliabilism, structurally parallel to consequentialism. In reliabilist theories the epistemic goal is to bring about true belief and to avoid bringing about bad states of affairs. And like most consequentialist ethics, reliabilism understands the good as quantitative. Whereas the utilitarian aims to maximize the balance of pleasure over pain, the reliabilist aims to maximize the balance of true over false beliefs.

(Zagzebski 1996, pp. 7-8)

If this picture of reliabilism is correct, then its structural similarity to consequentialism might generate some difficulty in accommodating the existence of ESAs.

Consequentialists famously have trouble countenancing MSAs. If we are obligated to perform that act that brings about the best consequences, then it is hard to see how there could be any act that is good to do but not morally required: either the act is the *best* one we could do and is thus required of us, or the act is *not* the best one we could perform and is thus impermissible. In either case, the act will not be morally supererogatory. Reliabilism may have a parallel

problem with countenancing ESAs. If reliabilism strives to maximize the balance of true beliefs over false ones, then the same structural dilemma arises. If the act is the one that maximizes true beliefs held and minimizes false beliefs held, then the act is epistemically required. Conversely, if the act fails to do this, then the act is epistemically impermissible. In neither case would the act be epistemically supererogatory.

The majority of reliabilists do not adhere to this maximizing view; rather, they tend to be satisficers about the process by which one acquires beliefs. There is a certain threshold of reliability that must be passed for a belief to be justified. This position is also consistent with reliabilism's parallels to consequentialism since not all forms of consequentialism are committed to maximizing the good. Satisficing consequentialism, for instance, claims that we just have to perform the act that is "good enough"—we do not always have to do what is best (Slote 1984). Similarly, progressive consequentialism requires us to improve the world (i.e., to leave the world better off than if we did nothing) but make it as good as we possibly can (Jamieson and Elliot 2009). Since both these theories allow for the existence of morally good actions that are not morally required, they appear able to accommodate MSAs. In a similar manner, satisficing reliabilism appears able to accommodate ESAs, so reliabilists may not have much to fear. However, matters may not be so simple. Satisficing consequentialism is not without difficulties of its own (see, e.g., Bradley 2006), and satisficing reliabilism may encounter similar difficulties. Further exploration of this matter seems crucial because if the most plausible form of reliabilism is a maximizing form, then the existence of ESAs would be a significant obstacle to this view of epistemic justification. Much like MSAs might give us some reason to think that approaching morality solely based on consequences is misguided, ESAs might give us some reason to favor

approaches to epistemic justification that do not focus exclusively on the reliability of truth-forming processes.²³

Though I have sketched only a few potential implications of epistemic supererogation, this sample should be enough to demonstrate just how significant ESAs are, should they indeed be possible. They may influence our understanding of contemporary epistemology a great deal, and the search for how best to incorporate them into accounts of epistemic duties is one that we ought to pursue in earnest.

7 Conclusion

Through the examination of five distinct cases, I have argued that the existence of epistemic supererogation is highly plausible because this classification provides the best explanation for our normative evaluations of certain epistemic acts. Without ESAs, we would have difficulties explaining our attitudes about a wide range of cases where conduct appears epistemically praiseworthy but (for various reasons) does not seem like it could be epistemically required. The fact that MSAs exist and that ethics and epistemology should be (and often are) conceived as isomorphic frameworks gives us further reason to think ESAs should be incorporated into our thinking about epistemic duties and epistemology more generally even though epistemic supererogation has gone largely unnoticed and unexplored up to this point. The significance of ESAs should not be underestimated: epistemic supererogation provides a new means of evaluating epistemic duties and may have further implications regarding theories of epistemic justification.

²³ The deontologically-oriented approaches to epistemic justification that Zagzebski mentions would be one way of accomplishing this task, but they may not be the only way of doing so. A virtue-based approach to epistemic justification, to give just one example, might also work.

Epistemic supererogation represents largely uncharted territory in the theoretical landscape, and further exploration of it seems to be one of the next important tasks facing contemporary epistemologists, particularly since my own analysis has focused exclusively on *actions* that are epistemically supererogatory. We can ask the related question of whether there might be *cognitive states* that are epistemically supererogatory, and if we get an affirmative answer, then epistemic supererogation may prove an even richer area for continued inquiry. Of course, there are some important obstacles to thinking such a category of cognitive states exists. One might think, for instance, that the lack of voluntary control over our cognitive states renders evaluations of praise and blame for them inappropriate.²⁴ We may be praiseworthy or blameworthy for actions that we can voluntarily perform, but can the same really be said about cognitive attitudes that are involuntarily adopted or only within our voluntary control to some minimal and indirect extent? Praiseworthiness would be a component of an epistemically supererogatory cognitive state, and so if praise and blame cannot be appropriately ascribed to cognitive states, then none of them will be epistemically supererogatory. Whether this obstacle (and others) can be overcome, however, is a discussion that must wait for another occasion. For the moment, we can conclude only that there are epistemically supererogatory *actions*, though that conclusion is far from insignificant.

Acknowledgments Acknowledgments removed for blind review.

References

- Alston, W. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2, 257-299.
- Attfield, R. (1979). Supererogation and double standards. *Mind*, 88, 481-499.

²⁴ I borrow this sentiment from Alston (1988).

- Bradley, B. (2006). Against satisficing consequentialism. *Utilitas*, 18, 97-108.
- Clifford, W. K. (1877). The ethics of belief. *The Contemporary Review*, 29, 289-309.
- Feldman, R. (2002). Epistemological duties. In P. Moser (ed.), *The Oxford Handbook of Epistemology*. New York: Oxford University Press.
- Firth, R. (1981). Epistemic merit, intrinsic and instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55(1), 5-23.
- Gilovich, T. (1991). *How We Know What Isn't So*. New York: The Free Press.
- Grimm, S. (2009). Epistemic normativity. In A. Haddock, A. Millar, and D. Pritchard (eds.), *Epistemic Value*. New York: Oxford University Press.
- Hale, S. (1991). Against supererogation. *American Philosophical Quarterly*, 28(4), 273-285.
- Hall, R. J., and Johnson, C. R. (1998). The epistemic duty to seek more evidence. *American Philosophical Quarterly*, 35(2), 129-139.
- Hsieh, N. (2007). Incommensurable values. *Stanford Encyclopedia of Philosophy*. Accessed April 28, 2014, from <http://plato.stanford.edu/>.
- Jackson, M. W. (1986). The nature of supererogation. *The Journal of Value Inquiry*, 20, 289-296.
- Jackson, M. W. (1988). Above and beyond the call of duty. *Journal of Social Philosophy*, 19(2), 3-12.
- Jamieson, D., and Elliot, R. (2009). Progressive consequentialism. *Philosophical Perspectives*, 23, 241-251.
- Kawall, J. (2005). Promising and supererogation. *Philosophia*, 32, 389-398.
- Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy*, 105(10), 611-633.

- Kim, K. (1994). The deontological conception of epistemic justification and doxastic voluntarism. *Analysis*, 54(4), 282-284.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, 92(1), 33-48.
- Lord, C., Ross, L., and Lepper, M. (1979). Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11), 2098-2109.
- McCarty, R. (1989). Limits of Kantian duty, and beyond. *American Philosophical Quarterly*, 26(1), 43-52.
- Mellema, G. (1991). *Beyond the Call of Duty: Supererogation, Obligation, and Offence*. Albany: SUNY Press.
- Plantinga, A. (1986). Epistemic justification. *Noûs*, 20(1), 3-18.
- Plantinga, A. (1988). Positive epistemic status and proper function. *Philosophical Perspectives*, 2, 1-50.
- Plantinga, A. (1991). The prospects for natural theology. *Philosophical Perspectives*, 5, 287-315.
- Postow, B. C. (2005). Supererogation again. *The Journal of Value Inquiry*, 39, 245-253.
- Pybus, E. (1982). Saints and heroes. *Philosophy*, 57(220), 193-199.
- Sikora, R. I. (1979). Utilitarianism, supererogation, and future generations. *Canadian Journal of Philosophy*, 9(3), 461-466.
- Slote, M. (1984). Satisficing consequentialism. *Proceedings of the Aristotelian Society*, 58, 139-163.
- Stanlick, N. (1999). The nature and value of supererogatory actions. *Journal of Social Philosophy*, 30(1), 209-222.

Tidman, P. (1996). Critical reflection: an alleged epistemic duty. *Analysis*, (56)4, 268-276.

Urmson, J. O. (1958). Saints and heroes. In A. I. Melden (ed.), *Essays in Moral Philosophy*.

Seattle: University of Washington Press.

Vessel, J. (2010). Supererogation for utilitarianism. *American Philosophical Quarterly*, 47(4),

299-319.

Zagzebski, L. (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical*

Foundations of Knowledge. Cambridge: Cambridge University Press.