



IN DEFENSE OF THE PRIMACY OF THE VIRTUES

BY JASON KAWALL

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 3, No. 2 | AUGUST 2009

URL: WWW.JESP.ORG

COPYRIGHT © JASON KAWALL 2009

In Defense of the Primacy of the Virtues

Jason Kawall

IN RECENT DECADES THERE HAS BEEN a great deal of interest in virtue ethics, broadly construed. There are, of course, many different such theories, and some dispute over what conditions a theory must meet to qualify as a virtue ethics.¹ In what follows, I respond to a set of basic objections to a paradigmatic set of such theories – those virtue ethics according to which moral properties such as rightness and goodness (and their corresponding concepts) are to be explained in terms of the virtues or virtuous agents (and similarly with their corresponding concepts).

The basic intuition underlying the objections is that our virtue concepts (or, indeed, the concept of a virtue, *tout court*) must be derivative from other, more fundamental moral concepts. Similarly, the virtues themselves – as well as their value – are thought to be best understood in terms of the right or the good. Virtue theorists have, in some way, confused cart and horse. Consequentialists will treat the virtues as character traits that serve to maximize (or produce sufficient quantities of) the good, where the good is taken as explanatorily basic. Deontologists will understand the virtues in terms of dispositions to respect and act in accordance with moral rules, or to perform morally right actions, where these moral rules or right actions are fundamental. Furthermore, the virtues will be considered valuable just insofar as they involve such tendencies to maximize the good or to perform right actions.

In contrast, the forms of virtue ethics that I wish to defend would satisfy the following four conditions:

- (i) The concepts of rightness and goodness would be explained in terms of virtue concepts (or the concept of a virtuous agent).
- (ii) Rightness and goodness would be explained in terms of the virtues or virtuous agents.
- (iii) The explanatory primacy of the virtues or virtuous agents (and virtue concepts) would reflect a metaphysical dependence of rightness and goodness upon the virtues or virtuous agents.
- (iv) The virtues or virtuous agents themselves – as well as their value – could (but need not) be explained in terms of further states, such as health, eudaimonia, etc., but where these further states do not require an appeal to rightness or goodness.²

¹ For a recent effort to provide a taxonomy, see Sean McAleer, “An Aristotelian Account of Virtue Ethics: An Essay in Moral Taxonomy,” *Pacific Philosophical Quarterly* 88 (2007), 208-225. See also Michael Slote, *Morals from Motives* (New York: Oxford University Press, 2001), and Gary Watson, “On the Primacy of Character” in Owen Flanagan and Amelie Rorty, eds., *Identity, Character, and Morality* (Cambridge, MA: MIT Press, 1990), 449-70.

² With the inclusion of this condition we capture those virtue ethics that Slote refers to as agent-prior (where the virtues or virtuous agents are explained in terms of further states, such as flourishing), in addition to those theories that Slote refers to as agent-based (where the virtues or virtuous agents are treated as basic, without further explanation). See Slote (2001), 3-10.

“Rightness” in the above can be taken as standing in for all deontic statuses (that is, for example, wrongness is also to be treated in terms of the virtues or perhaps the vices); similarly, “goodness” can be taken as standing in for badness, and other such axiological statuses. One qualification is in order with respect to states such as eudaimonia or health. These might appear to be instances of goodness; however here we can draw a distinction between the goodness of a kind (particularly as a good agent or good creature of some kind) as reflected in health or flourishing (one is a healthy *human*, or is leading a good *human* life), versus goodness *tout court*. Only the latter instances of goodness are to be taken to be dependent upon the virtuous or virtuous agents for present purposes.³

It is worth stressing that not all theories that could plausibly be understood as forms of virtue ethics would satisfy the above conditions; the current goal is not to defend all possible virtue ethics. Rather, we are examining what might be taken to be among the more radical possible forms of virtue ethics, particularly in treating the virtues as explanatorily prior both to rightness and to goodness *tout court*. Why focus on these more radical forms? First, several prominent virtue ethics can be understood as satisfying the above conditions, including those of Michael Slote, Linda Zagzebski, and, perhaps (if controversially), Aristotle’s paradigmatic virtue ethics.⁴ Beyond this, many of the arguments presented here could be taken on board by those defending more moderate forms of virtue ethics, such as Rosalind Hursthouse or Christine Swanton (against those who would attempt to argue for the explanatory primacy of the right or of the good, for example).⁵ Thus the range of interest for most of these arguments will extend beyond those focusing on the more radical approaches.

The paper consists of four main sections, each dealing in depth with particular variations of the underlying worry that our virtue concepts must be derivative from other, more fundamental moral concepts (and similarly with the virtues themselves and other moral properties):

³ Thus the recent virtue consequentialisms of Julia Driver and others would not satisfy these conditions, insofar as they treat the virtues in terms of producing the good (where this presumably includes states other than goods of a kind such as health or flourishing). See Julia Driver, *Uneasy Virtue* (New York: Cambridge University Press, 2001).

⁴ Michael Slote (2001), and “Sentimentalist Virtue and Moral Judgement: Outline of a Project,” *Metaphilosophy* 34 (2003), 131-43. See also Linda Zagzebski, *Divine Motivation Theory* (New York: Cambridge University Press, 2004). McAleer (2007) provides an interpretation of Aristotle that would be congenial to this claim.

⁵ Hursthouse explicitly rejects attempts to “reduce or derive” all moral concepts to / from virtue concepts (or any others, including the right or the good); see her “Are Virtues the Proper Starting Point for Morality?” in James Dreier, ed., *Contemporary Debates in Moral Theory* (Malden, MA: Blackwell, 2006), 99-112; see in particular 100-1. For a more complete exposition of her approach, see Rosalind Hursthouse, *On Virtue Ethics* (New York: Oxford University Press, 1999). For target-centered accounts of the virtues, see Christine Swanton, *Virtue Ethics: A Pluralistic View* (New York: Oxford University Press, 2003), and Ronald Sandler, *Character and Environment: A Virtue-Oriented Approach to Environmental Ethics* (New York: Columbia University Press, 2007).

- (i) That the virtues can be properly explained as simple dispositions to perform right actions or promote the good (as reflected in our ability to pick out agents as virtuous by the acts they perform), where rightness or goodness are thus explanatorily primary;
- (ii) That we can identify certain actions as morally *right* without appeal to virtuous agents – suggesting a primacy of the right, where right actions are not properly understood in terms of the virtues or virtuous agents (contrary to claims that the virtues would be explanatorily primary);
- (iii) That we can identify certain states of affairs as *good* (or not), without appeal to virtuous agents and, furthermore, that appeal to virtuous agents to explain such value seems implausible (and not merely redundant) – suggesting a primacy of the good; and
- (iv) That virtue ethics face a dilemma – either virtuous agents are consistent in their actions and valuations – in which case the best explanation of this is that they are reacting to antecedently right actions or good states of affairs (and thus the virtues or virtuous agents would not be explanatorily primary), or they are not guided by prior right actions or good states of affairs. If so, the virtuous would be arbitrary and potentially inconsistent in their actions and valuations (thus rendering virtue ethics implausible).

In the concluding section I acknowledge final, broader issues that will ultimately need to be addressed in arriving at a viable virtue ethics (of the radical kind defended here). I do not attempt to fully answer these final questions in this paper; my goal is instead to address the prominent, fundamental objections to virtue ethics discussed in sections I – IV.

I

A simple version of the first objection is grounded in moral epistemology. When we identify honest people, benevolent people, and so on, we typically pick them out by the very fact that they perform honest or benevolent actions. Thus, we identify virtuous persons (or their possession of given virtues) through their performance of certain kinds of actions; as such, these right actions are explanatorily basic or primary, and the virtues are derivative and best understood as dispositions to perform these actions. Or so the objection would run.

Still, even if it were true that we identified the possession of various virtues primarily through agents performing certain characteristic kinds of actions, this would not demonstrate that benevolent actions are explanatorily prior to benevolence. To see this, compare the case of a disease and its symptoms. Clearly, we often identify the presence of a disease through its symptoms, yet we would not claim that these symptoms cause or explain the presence of the disease. And when we examine our concepts of disease and symptom, we best explain the notion of a symptom in terms of the more basic concept of a disease – even if we have better epistemic access to the presence of actual symptoms than to the presence of the more fundamental disease.

Similarly, in the case of the virtues, while we might use certain kinds of actions to pick out those who possess the virtues, we can readily hold that these sorts of actions are seen as honest or benevolent simply insofar as they are actions that would be characteristic of virtuous agents. Put another way, we can use virtuous actions as a heuristic in identifying virtuous agents, but we can use these actions in this way precisely because they are the sorts of actions that virtuous agents would perform. Furthermore, and crucially, we must distinguish between how we identify those who possess the virtues, and how we are to understand the relationships between rightness and virtue (and also the relationships among our various moral concepts). The epistemic methods we use to identify virtuous agents do not necessarily reflect the structure of our moral conceptual framework or the most plausible explanatory relationships between virtues and actions or states of affairs.

The objection could be modified now to the claim that virtues could not ultimately be anything more than dispositions to perform certain sorts of actions. If so, the analogy with diseases and symptoms would break down, as (unlike a disease that is more complex than a mere disposition to produce certain symptoms) there would be nothing more to a virtue than its production of right actions. We might call this “reductive deontologism.” One could insist that benevolence simply is a disposition to perform actions which improve the well-being of others (and perhaps oneself); there would be no “remainder” to benevolence beyond this.⁶ And similar claims could be made concerning other virtues – honesty, justice, and so on.

In response, consider a case in which a person loves another, perhaps a romantic love. And now consider how we should understand the actions of this person, or those who love in this way. An understanding analogous to that of reductive deontologists with respect to the virtues would have us hold that loving people are simply those who are disposed to perform independently grounded loving actions. To be in love would be a matter of performing (or being disposed to perform, under appropriate circumstances) such actions as sending flowers or what-have-you. Or perhaps there could be rules of love, such that being in love is simply a matter of acting in accordance with these rules.... But surely this sort of approach cannot be correct. Rather, we characterize actions as loving precisely because they are the sorts of actions that people in love might commonly perform. And notice that while we could well come to believe that a given person is in love with another on the basis of her actions, this does nothing to show that we ought to understand her being in love as merely a disposition to perform certain sorts of actions. Loving actions are the actions that flow out of the rich emotional

⁶ Or, if there were any such remainder, it would not be sufficient to allow the virtues to be explanatorily basic.

state of being in love; love is not a mere disposition to perform particular actions or follow a set of rules.⁷

We have then a case where it is clearly mistaken to understand a complex inner state as a mere disposition to perform certain independently identified actions. Virtue theorists⁸ can hold that we should understand the virtues in a similar fashion. Honest people are not merely disposed to perform independently grounded honest actions, any more than loving people are merely disposed to perform independently grounded loving actions. We could perhaps develop this point further by understanding honesty as a love of truth, benevolence as a general love of others (perhaps a form of agape), and so on. Of course, this is merely suggestive, and an analysis of the loves involved would be essential.⁹

But our main point remains – there is no reason to hold that virtue theorists must be placing the cart before the horse in holding that the virtues are explanatorily basic to their corresponding actions. The relationship between loving and loving actions can serve, at the very least, as a model for the virtue theorists’ understanding of the relationship between the virtues and virtuous actions. This is not yet to show that they are correct in appealing to this model; but my goal here is simply to remove initial, intuitive objections that suggest that any pure virtue theory *must* be mistaken.

We can also appeal to recent work by Christine Swanton in rejecting construals of the virtues as mere dispositions to perform right actions.¹⁰ Swanton argues that virtues, properly understood, involve a range of “modes of moral responsiveness.” She writes:

[The virtues] recognize that we are not only agents who are active in changing the world by promoting good (often at the expense of causing harm), but also agents who love and respect (often at the expense of maximizing good). Finally, they accept that we are not only active beings hell-bent on change, but are also passive in a sense: in our openness to, receptivity to, and appreciation of value and things.¹¹

⁷ A parallel reductive consequentialist understanding might hold that a loving person is simply disposed to perform actions that maximize the well-being of the beloved. But of course much more would need to be involved – a desire to be with the beloved, taking pleasure in the happiness of the beloved, (typically) a sexual attraction to the beloved, and so on. There would be a wide range of affective, conative, and other dispositions involved; loving a person cannot be adequately explained as a mere disposition to perform actions that maximize this person’s well-being.

⁸ Here, and throughout the body of this paper, “virtue theory” and “virtue theorist” should be understood in terms of those theories that meet the relatively demanding criteria noted in the introduction. Again, while there are other virtue theories that will not meet these criteria, they will not be our focus here.

⁹ I do not mean to endorse this proposal (of treating virtues as all involving some form of love); it is simply illustrative of how one might develop an account of the virtues.

¹⁰ Swanton (2003).

¹¹ Swanton (2003), 23.

Consider: the virtue of honesty does not consist in a disposition to tell as many truths as possible to as many people as possible.¹² Drawing upon Swanton's work, we could understand honest actions as those which would be performed (or approved of) by agents who possess the virtue of honesty, where this involves an appreciation or love of truth, a respect for other agents as autonomous, and so on.¹³ A person who values truth and respects others (in addition to other modes of moral responsiveness) will naturally tend to approve of various actions (in addition to being disposed to act in various ways); we need not assume that they are simply inclined to perform actions independently identified as "honest."

Still, we might now press a final variation of the worry that the virtues can be reduced to dispositions to perform right actions or to maximize the good. Not all deontologists and consequentialists would attempt to reduce the virtues to *mere* dispositions to perform certain actions or to maximize the good. They could instead agree that the virtues are much richer dispositions, including the range of attitudes and modes of responsiveness suggested by Swanton above. But they could also hold that the *value* of these broader, richer virtues is still to be explained in terms of their tendency to issue in right actions or good states of affairs, where these latter are treated as foundational or basic. As such, the value of the virtues would still be derivative from (and best explained in terms of) antecedently valuable states or actions.

Important recent work by Thomas Hurka can be seen as providing a basis for the objection, one where the value of the virtues would be best explained in terms of the good.¹⁴ Hurka provides a recursive account of the virtues, starting with the claim that there are certain basic intrinsic goods:

(BG) Pleasure, knowledge, achievement are intrinsically good.¹⁵

The recursive aspect enters as Hurka suggests that certain attitudes toward such basic goods themselves are intrinsically good:

(LG) If x is intrinsically good, loving x (desiring, pursuing, or taking pleasure in x) for itself is also intrinsically good.¹⁶

¹² Swanton herself treats honesty as a mere disposition to tell the truth (or at least not lie) and, as such, does not consider it to be a genuine virtue. On the other hand, she suggests that there is a distinct virtue "of a correct disposition with respect to divulging information." It seems most natural to me to maintain that honesty is this latter virtue, rather than a mere disposition of truth-telling. See Swanton (2003), 246.

¹³ And it is worth noting that respecting a person primarily involves certain attitudes and a valuing of the person herself; respect is not plausibly construed as a mere disposition to perform "respectful" actions. Again, it seems best to understand respectful actions as those that would be performed (or approved of) by a person with an attitude of respect (toward a given person or object).

¹⁴ Thomas Hurka, *Virtue, Vice, and Value* (New York: Oxford University Press, 2001).

¹⁵ Hurka (2001), 12. Note that Hurka is not committed to this particular list of basic goods; he simply takes these to be plausible candidates.

¹⁶ Hurka (2001), 13.

We can pass over certain features of Hurka's account – he suggests a set of basic evils, and holds that loving what is evil is itself evil, while hating what is good is evil, and so on. The intuitive structure should be apparent. With this in hand, Hurka proposes the following definitions of moral virtue and vice:

The moral virtues are those attitudes to goods and evils that are intrinsically good, and the moral vices are those attitudes to good and evils that are intrinsically evil.¹⁷

Focusing on the virtues, we can now see how the objection might be sustained. Hurka can allow that the virtues involve various modes of responsiveness to intrinsic goods (pursuing, taking pleasure in, etc.) while also holding that the value of these responses is explained by, and derivative from, the value of the basic goods. As he says, “In my view, virtue cannot plausibly be treated as morally primitive: it consists in an appropriate response to other moral considerations and must be analyzed in terms of those considerations.”¹⁸

A response to this proposal will emerge and be defended over the next two sections. But to anticipate somewhat: I will argue that the virtue theorist need not accept that there are basic intrinsic goods (pleasure, knowledge, etc.) or evils (pain, ignorance, etc.) that are prior to the attitudes of the virtuous. That is, while there are obviously pleasures and pains, we can plausibly deny that these have a status of being independently, intrinsically good or evil. In turn, this will allow the possibility of understanding even these goods and evils in terms of the responses of the virtuous.

II

We can now turn to our second objection to the primacy of the virtues; in so doing, we will also arrive at a response to the final variant of the first. The first objection (and its variants) focused on the proposal that the virtues are, by their very nature, simply dispositions to follow rules or perform certain actions (as reflected in our common practice of attributing character traits to agents based upon their actions) – or that the value of the virtues can be understood simply in terms of their relationship to right actions or good states of affairs. The second objection draws upon the fact that we often pick out paradigmatic right (and wrong) actions or rules in our moral thinking without any appeal to virtuous people and their behavior. This fact about moral epistemology suggests that the rightness of these actions or rules is not derivative from the virtues in any way – after all, a good explanation of this epistemic fact would be that the right is explanatorily prior to the virtues (or the virtuous), and that this reflects a metaphysical dependence of virtue upon the

¹⁷ Hurka (2001), 20.

¹⁸ Hurka (2001), 42.

right. We could then understand virtuous people to be those who are disposed to perform these actions or act in accordance with these rules.¹⁹ We would thus embrace the primacy of the right over the virtuous.

To bring this out more clearly, consider the following examples. We can see that torturing puppies and children for pleasure is wrong directly; there is no need to appeal to how virtuous people would act. Similarly with acts of murder, rape, and so on; the actions are clearly wrong, and any appeal to the attitudes or behaviors of the virtuous in such cases seems superfluous at best. Upon reflection, we see that a rule forbidding lying (*ceteris paribus*) seems correct; again, there is no need to appeal to virtuous agents.

The following passage from Jonathan Dancy (cited by Simon Kirchin) can be understood as capturing the underlying intuition:

[T]aking our experience at face value, we judge it to be experience of the moral properties of actions and agents in the world. And if we are to work with the presumption that the world is the way our experience represents it to us as being, we should take it in the absence of contrary considerations that actions and agents do have the sorts of moral properties we experience in them. This is an argument about the nature of moral experience, which moves from that nature to the probable nature of the world.²⁰

Focusing on actions, the claim is that we seem to experience them as having genuine, independent moral properties and, assuming that our experience of the world is generally accurate, this suggests that there are such properties of rightness, etc., in the world prior to our experiences. The virtues (or virtuous characters) cannot be explanatorily basic in such cases; rather, these actions are antecedently and independently wrong. Nor is the point merely epistemic – it would be wrong-headed to explain the wrongness of these actions by appeal to the virtuous. Torturing children is not wrong because good people would not do such things; it is wrong because, at the very least, it causes suffering in an innocent. And we could understand virtuous people as those

¹⁹ The first objection is based on the claim that the essence of virtues must lie in being dispositions to perform certain actions (including, perhaps, maximizing the good). The second objection is based on the claim that we can simply see that a wide range of actions are right (or wrong), without any appeal to virtues; as such, the virtues cannot be explanatorily basic in these cases.

²⁰ Jonathan Dancy, “Two Conceptions of Moral Realism,” *Proceedings of the Aristotelian Society* supp. vol. 60 (1986), 172. Dancy’s paper is cited in Simon Kirchin, “Ethical Phenomenology and Metaethics,” *Ethical Theory and Moral Practice* 6 (2003), 241-64, and is reprinted in James Rachels (ed.), *Ethical Theory* (New York: Oxford University Press, 1998). More broadly, there has been a resurgence of interest in ethical intuitionism, a view that (if correct) could undermine a pure virtue ethics. Philip Stratton-Lake notes that “Intuitionists hold that certain basic moral propositions are self-evident, and thus can be directly known by intuition.” On such views, the truth of such moral propositions is independent and prior to the judgments of the virtuous. See Philip Stratton-Lake (ed.), *Ethical Intuitionism: Re-evaluations* (New York: Oxford University Press, 2002). [Quotation from Stratton-Lake’s introduction to the volume, 18.]

who are disposed to act in accordance with these antecedently identified rules or actions.

One line of response available to virtue ethicists would draw attention to the fact that when we consider such cases more closely, we find that we often need to appeal to the motives and character of an agent in order to fully determine the moral status of the actions. For example, in the case of torturing children, we cannot explain the wrongness simply in terms of causing suffering in an innocent – after all, surgery might cause such suffering, too. The motives and intentions of an agent can be relevant to assessments of rightness and wrongness (even in apparently obvious cases). Similarly with our earlier case of giving someone flowers – if done out of certain appropriate motives, the act is loving and presumably right. But if the action is intended to set off the recipient’s allergies, the action is no longer loving (and, indeed, could be quite cruel or hateful). It seems that we often cannot fully characterize actions as right (or wrong, or praiseworthy, or loving, etc.) independently of an agent’s motives. If so, we would thus avoid a pure primacy of the right: instead, characterizing actions as right would (often) require appeal to such things as virtuous motives and intentions.²¹

More broadly, the virtue theorist can grant that we seem to simply see some actions as right “directly,” without complicated reasoning about what virtuous people would do. But, crucially, when we do our moral thinking or respond morally to situations, we do not do so as blank slates. Ordinary people possess various virtues to at least some degree. We may see torturing puppies as wrong immediately, but this is because we possess a basic compassion, sense of justice, and so forth. We see promising as providing a prima facie duty, but this is because we respect others and value the truth. Our responses depend, of course, on the kind of creature that we are and the kind of character that we possess. Perhaps there are Martians or robots who would not respond in these ways when presented with tortured puppies. It is far from clear that when we (basically decent, ordinary) humans respond to certain actions as wrong (or right) immediately, that we are responding to some prior, independent rightness to the actions themselves. Rather, it is

²¹ This proposal would also allow a virtue ethicist to embrace an intuitionism of the kind considered in the previous footnote, if she were to hold that any intuited moral propositions would require reference to virtuous motives or intentions. For further cases where the motives of actions are crucial to determining their rightness, see Michael Stocker’s classic “The Schizophrenia of Modern Ethical Theories,” *Journal of Philosophy* 73 (1976), 453-66. While the current line of response would be appealing to many virtue ethicists, I personally worry that an appeal to motives will not always be required in determining the rightness of an action, and that (similarly) we cannot determine that an action is wrong simply because it is performed out of vicious motives. For example, if Claire saves a drowning person from a river solely in order to receive a reward, it seems plausible to hold that she performs a right action, but with a vicious motive. I do not wish to take a stand here on whether appeals to an agent’s motives are always required in determining the rightness (or wrongness) of her actions, or whether bad motives thereby entail wrong actions. I will, however, next present an alternative response available to virtue ethicists who are unsure of the current response.

quite plausible to hold that our responses are the result of training, natural tendencies to sympathy, anticipating the responses of those around us, etc.²² And, if this is correct, there is no difficulty in maintaining that the virtues are explanatorily basic, even in these cases.

A further analogy will help to clarify the current proposal. Consider a normal, adult human living in a contemporary Western society, looking at a flag. Such a person will immediately see the Union Jack *as* a Union Jack, a symbol of the U.K., and will similarly see the American stars and stripes as a symbol of the United States. Indeed, it would be difficult for such a person to see an American flag as merely a series of colored lines and patterns rather than as a flag and symbol. But we have been trained about flags as symbols of nations – obviously a Martian or a human without such experience will not respond with immediate thoughts of the U.K. or the U.S. upon being presented with the same flags; indeed, presumably they will not even see them *qua* flags. And notice how implausible it would be to posit an independent, prior “Britishness” that inheres in the Union Jack pattern, such that we are simply reacting to this innate Britishness when we see the flag immediately as a symbol of the U.K. The current suggestion, then, is that, in a similar way, normal adults might “automatically” see the wrongness of certain actions, but that this is a result of training (combined perhaps with certain natural / evolved dispositions) – it need not be treated as a response to a prior, independent property of wrongness. And when philosophers react to thought experiments in ethics (and “immediately” perceive some actions to be wrong, others right), we do not do so as blank slates; we have developed character traits that shape our reactions. We thus do not need to posit antecedently right actions even in those cases where we seem to simply see certain actions as right; the primacy of the virtues is unaffected.

While this section has focused on responding to claims of the explanatory primacy of the right, lessons can be drawn for claims concerning the primacy of the good. Returning to Hurka and (the final variant of) the first objection, we can grant that it may similarly seem immediately obvious to us that pleasure, knowledge, and achievement are good, while pain and the like are intrinsically bad. But just as in the case of the obviously right and wrong actions considered in this section, we can plausibly deny that these states have their value-status independently of, or prior to, the judgments of the virtuous. Seeing pleasure as good need not involve perceiving an independent, prior value, any more than seeing the Union Jack as the flag of the U.K. needs to involve perceiving an independent, prior “Britishness” in the pattern. We are decent people with evolved dispositions and training that lead us to such assessments. In the following section we will consider objections proposing the primacy of the good over the virtuous in greater depth.

²² Kirchin provides additional responses against this sort of appeal to ethical phenomenology as supporting a form of mind-independent moral realism. See Kirchin (2003), 248-51.

III

The next objection to the virtue theorist grants that our intuitions and reactions to various cases might be a result of our character and training, rather than being responses to independent rightness or wrongness in the world. But the question is pressed: why do virtuous people respond in these ways to these cases? Why do normal or good people respond negatively to acts of puppy-torturing or murder? Consider the following from a recent review essay by David Copp and David Sobel:

It is facts about the alternatives a person must decide among, including such things as the impact the alternatives will have on people's ability to meet their needs, that determine what a person ought to do. It is not facts about what a virtuous person would want her to do, or facts about the motives that the person would actually be acting from if she were to do the various alternatives. If someone is drowning, for example, and if you can save her at no risk and at negligible cost to yourself, you ought to save her because otherwise her life will be wasted. It is because a life would otherwise be wasted that a virtuous person would want you to save her.²³

Surely helping a drowning person is right because it prevents the suffering and potential death of the victim and the loss of future goods for the person, her friends, and so on; it is not right simply because virtuous agents would save the person. One could thus plausibly argue that the virtuous are simply responding to states of affairs with intrinsic value (or disvalue), and that it is this intrinsic goodness (or badness) that is explanatorily basic. This, of course, brings to mind Hurka's proposal concerning basic intrinsic goods.

But a virtue theorist can hold that concerns with various states of affairs are essential components of the virtues. On the current approach, the virtue theorist can agree with Copp and Sobel that virtuous agents would concern themselves with the potential suffering and loss of life of the drowning victim (etc.); all that is being denied is that there are prior properties of goodness or badness that the virtuous are responding to. That is, the virtuous are responding to the potential suffering and loss of life itself; there is not a prior property of "badness" attaching to the potential suffering and loss of life that causes the virtuous to react as they do.

For example, we notice that any unimpaired virtuous agent will be moved by suffering (and respond negatively to it), and will judge a state of affairs which includes suffering to be worse than a similar state of affairs without suffering.²⁴ This leads us to an alternative way of understanding the moral status of suffering. Gary Watson suggests that

²³ David Copp and David Sobel, "Morality and Virtue: An Assessment of Some Recent Work in Virtue Ethics," *Ethics* 114 (2004), 514-54. Quotation from 552.

²⁴ At least in standard cases – we can put aside circumstances of punishment and so on for present purposes.

To be sure, a concern for outcomes will be internal to certain virtues. For instance, the benevolent person will be concerned that others fare well. But the moral significance of this concern stems from the fact that it is part of a virtue, not from the fact that misery and well-being are intrinsically or ultimately bad and good respectively. To put it another way, it will follow from an ethics of virtue that virtuous people care about things (and outcomes) for their own sakes (as final ends in themselves). There is no further commitment, however, to the idea that these concerns are virtuous ones because their objects are inherently valuable or desirable for their own sakes.²⁵

Thus the virtues will often lead us to value certain outcomes but, according to the virtue theorist, the moral status of these outcomes is determined by the reactions and attitudes of the virtuous. The virtuous are responding directly to states of affairs – but there are not prior moral properties supervening on these states of affairs that are guiding the reactions of the virtuous. We thus have a plausible alternative to views like that of Hurka (that posit basic, independent intrinsic goods and evils in the world).

Now it might seem again that the virtue theorist has placed the cart before the horse. Surely (for example) the pain of the puppies being tortured is fundamentally what makes such an action wrong, and an appeal to the attitudes of virtuous agents misses the point. But note: the virtue theorist can entirely agree that virtuous agents respond negatively to torturing puppies *precisely because of the pain that would be caused*, and that, as such, *it is because of the pain that would be caused that the action would be wrong*. All that is being denied is the presence of any prior, independent goodness or badness (normative properties) that cause the reactions of the virtuous. We need to be careful to distinguish different senses of (and answers to) the question “Why is this action wrong?”.

To illustrate, consider a divine command theorist who holds that actions are right if and only if they are in accordance with God’s commands. Furthermore, suppose that God commands us to act as act utilitarians. Finally, imagine a person contemplating the specific action of torturing puppies. Why, on this account, would it be wrong to torture puppies? We can distinguish three answers, corresponding to three senses of the question:

Meta-ethical: It is wrong because it is contrary to God’s commands, and actions possess the property of wrongness insofar as they are not in accordance with God’s commands.

Normative: It is wrong to torture puppies because doing so does not maximize utility (and thus is wrong according to the standards of act utilitarianism).

Instantiation: It is wrong to torture puppies because it causes them terrible suffering.

Notice then, that on this theory, one could say both that it is wrong to torture puppies because it is contrary to God’s commands, and that it is wrong

²⁵ Gary Watson, “On the Primacy of Character” in Owen Flanagan (ed), *Identity, Character, and Morality: Essays in Moral Psychology* (Cambridge: MIT Press, 1990), 449-69.

to torture puppies because it causes them terrible suffering. The former answers a meta-ethical question: *what is it that makes actions (including this one) wrong?* The latter answers the question of *why is this particular action wrong?*

Or consider a different theory, a form of moral realism, according to which badness is a natural property that supervenes on various natural base properties (such as suffering), while goodness is a natural property that supervenes on other base properties (such as happiness). Assume further that rightness supervenes on actions that maximize utility (where we can treat this simply as a matter of the net balance of goodness over badness produced); we might add a property of wrongness that supervenes on all other actions.²⁶ We essentially have a moral realist form of act utilitarianism. Now consider again the torturing of puppies. Why would it be wrong to do so?

Meta-ethical: It is wrong to torture puppies because the property of rightness would not supervene on the action of torturing puppies (presumably because a great deal of badness would supervene on their suffering), while the property of wrongness would so supervene.

Normative: It is wrong to torture puppies because doing so does not maximize utility.

Instantiation: It is wrong to torture puppies because it causes them terrible suffering.

Again, we need to distinguish various senses of our question: the torturing of puppies is wrong because the natural, supervenient property of rightness fails to supervene on the action while the property of wrongness does so supervene on the action; it is also wrong because it fails to maximize utility or the net balance of goodness over badness; and it is also wrong because of the puppies' suffering itself.²⁷ Each of these claims is equally correct and corresponds to a different sense of the question "why is this action wrong?"

²⁶ Why do these moral properties supervene on the base properties that they do? The realist need not provide an explanation here, anymore than she needs to explain why the laws of physics or biology are as they are. This would simply be part of the nature of the world we live in. See Russ Shafer-Landau, *Moral Realism: A Defence* (New York: Oxford University Press, 2003), especially chapter 2.

²⁷ We can thus distinguish three senses in which one might claim that causing suffering is a wrong-making property. At the metaethical level this could be taken, as in the current example, as a naturalistic realism where wrongness supervenes on causing suffering. At the normative level, the claim can be understood as stating that causing suffering is contrary to a given set of norms. Thus causing suffering will be a wrong-making property on most plausible normative theories. (And notice how a utilitarian emotivist could embrace the claim that causing suffering is a wrong-making property in this normative sense, even while denying the above metaethical claim; for the emotivist, the wrongness of causing suffering metaethically will be understood in terms of the attitudes of agents toward causing suffering). Finally, to hold that causing suffering is a wrong-making property in an instantive sense would be to hold that what explains the wrongness of some particular action, why it runs contrary to the norms at stake, and why it possesses the property of wrongness (understood as metaethical claim), is that this particular action involves causing suffering.

So what should the virtue theorist say here? I would suggest the following:

Metaethical: It is wrong to torture puppies because virtuous agents would not perform such actions (or would disapprove of them in some fashion), and such actions would thus possess the property of moral wrongness.

Normative: It is wrong to torture puppies because virtuous agents [given their psychologies] would not perform or approve of such actions. These approvals and disapprovals do not follow a simple normative code like act utilitarianism.

Instantiation: It is wrong to torture puppies because it causes them terrible suffering (which will lead to disapproval by virtuous agents due to their compassion and benevolence).

The distinction between the responses to the metaethical and normative questions requires some elaboration. Intuitively, on the current proposal, virtue theories are metaethical rivals to various moral realisms (where these basic realist properties are goodness or rightness), contractualisms, and other non-realist positions. But they are also normative theories and, at the normative level, they are rivals to consequentialist and deontological theories. Compare a slightly modified version of the theory: suppose that virtuous agents would approve of or endorse act utilitarianism.²⁸ If this were so, we would keep the same metaethical answer, but our normative answer would now be that torturing puppies is wrong as it does not maximize utility. On the actual proposed account, it is held that the approvals and behaviors of virtuous agents do not follow a simple set of normative rules; as such, we appeal directly to the approvals or behaviors of virtuous agents at the normative level.

The instantiation question can be understood as asking why a particular action is such that it fails to meet the criteria of rightness given at the normative level. And here the virtue theorist's answer is that it is the suffering of the puppies that makes torturing them wrong. Virtuous people will be moved by such suffering; if the suffering were not present, they would not have the reactions that they do. The virtuous, given their particular psychologies, will respond negatively to actions because of such features as causing unnecessary suffering, being a case of lying, and so on. It is precisely because of the puppies' suffering that virtuous agents will deem torturing them to be wrong; in this (instantiation) sense, torturing puppies is wrong because it causes them to suffer. We thus need to clearly distinguish between why the virtuous will have the responses that they have, and the nature of moral properties. The virtue theorist can at once say both that actions of harming puppies are wrong because of the suffering caused (this will be what concerns the virtuous and causes them not to approve of such actions), and because the virtuous would respond negatively to the action (where this latter can be under-

²⁸ We might imagine here (as with our divine command example) that the psychology of these virtuous agents (or God) is such that they are so deeply moved by suffering and happiness that they endorse a form of act utilitarianism.

stood as both a normative or a metaethical claim).

With this in hand, we can see more clearly the error in the objection. Recall Copp and Sobel's statement:

It is facts about the alternatives a person must decide among, including such things as the impact the alternatives will have on people's ability to meet their needs, that determine what a person ought to do. It is not facts about what a virtuous person would want her to do, or facts about the motives that the person would actually be acting from if she were to do the various alternatives.

Copp and Sobel's objection confuses the instantiation and normative (and metaethical) senses of the question "Why is this action right / wrong?" The virtue theorist can entirely agree that, in one fundamental sense, it is precisely preventing the potential loss of life and suffering that would make saving a drowning person right. To see the flaw in this sort of objection, compare parallel objections to normative and metaethical rivals to virtue theories:

Utilitarianism: "But don't you see – torturing puppies is not wrong because it fails to maximize happiness or preference-satisfaction or anything abstract like that! It is wrong because the puppies would suffer." Notice that a utilitarian could properly hold that torturing puppies is wrong at the instantiation level precisely because of the suffering caused – and that it is because of this suffering that such actions fail to maximize happiness. The utilitarian is not overlooking the puppies' suffering.

Kantianism: "Killing a person is not wrong because doing so fails to satisfy the Categorical Imperative! It is wrong because you are ending a person's life and ending all of her possibilities, etc." Notice that a Kantian could properly hold that killing a person is wrong because of lost opportunities, suffering, and so on (at the instantiation level). And insofar as this is the case, such actions will not satisfy the Categorical Imperative (they will not reflect universalizable maxims, or they will involve treating the victim as a means only, etc.).

Supervenient Non-Natural Properties (a metaethical rival): "Torturing puppies is not wrong because there is some mysterious non-natural moral property that supervenes upon such actions! It is wrong because of the suffering of the puppies!" Again, notice that someone could hold both that the wrongness of the action simply is a matter of its possessing a supervening, non-natural property (metaethically), *and* that it is precisely because of the suffering of the puppies that "wrongness" would supervene on actions of puppy-torturing. The objection confuses the instantiation and metaethical senses in which an action is wrong.

Recall that (on the present account) the virtue theorists' claim is that it is the reactions of the virtuous that give actions, events, etc., their moral status. Briefly put: humans are creatures who have evolved over time with a range of tastes, dispositions, common preferences, and so on. It could be a mere descriptive fact about most humans that they care about the suffering involved in puppy-torturing – the mere result of social training and certain common human emotional dispositions. More broadly, while we find pains to be unpleasant, and have an aversion to them, this can all be understood descriptively: the mere fact that we are averse to pains does yet justify holding that pains must be intrinsically bad. There is a gap between the descriptive and

the normative here. On the current proposal, what changes things – what turns a mere disapproval common among many humans – into *moral wrongness* (or *badness*) is the fact that particular kinds of agents under particular kinds of conditions would share a disapproval of this sort. It is the fact that *virtuous* agents in particular, (with their particular psychologies and dispositions), would have negative reactions that makes such acts *morally wrong* and the pains *morally bad* (at the metaethical and normative levels) – rather than merely displeasing to most humans. But at the same time, there is no need to downplay the importance of the suffering involved in puppy-torturings (or our own pains) on the virtue theorists’ view; indeed this suffering will be the focus of the concerns of the virtuous (and the reason why the action is wrong, in the instantiation sense). The only claim is that these actions, states of affairs, and so forth take on a *moral* status insofar as virtuous individuals in particular will have certain attitudes toward them. There are not prior, independent moral facts toward which the virtuous merely respond.

We can turn to a final variation of this worry: it could be objected that even if virtuous agents are not responding to antecedent rightness, intrinsic goodness or what-have-you, they are simply responding appropriately to relevant features of given situations. That is, even if there is not a prior property of wrongness that the virtuous are responding to, it is the suffering of tortured puppies itself (for example) that makes the attitudes of the virtuous (disapproval) appropriate or justified. There are features of situations which make certain attitudes appropriate – it is such features that give the attitudes of the virtuous their status as morally appropriate, rather than the attitudes of the virtuous determining the moral status of these features themselves. As such, the virtuous and their attitudes would not be explanatorily basic.

While this objection has a certain initial appeal, we need to bear in mind the discussion of section II. True, it seems that suffering warrants disapproval or hatred – and this might strike us as an immediate judgment. But again, we “ordinary folk” typically have at least somewhat virtuous characters, and have been trained to react in certain ways – so we tend to think that the suffering of others makes disapproval appropriate. There is no need to hold that it is an antecedent, independent fact that pleasure warrants or justifies approval. Rather, we can plausibly hold that through training and “natural” human tendencies, we will come to deem such things as pain, pleasure, knowledge, and achievement as warranting certain attitudes on our part.

IV

The next objection can be understood as a dilemma for the virtue theorist that arises if we grant that neither the good nor the right are explanatorily prior to the virtuous. It is perhaps worth noting that similar objections can be raised against ideal observer theories, and that they bear a striking resemblance to the Euthyphro dilemma commonly raised against divine command theories. Indeed, the objection is perhaps best understood as a form of the

Euthyphro dilemma applied to virtue ethics. Peter Singer provides a clear statement of the dilemma, as posed for the divine command theorist:

If all values result from God's will, what reason could God have for willing what he does? If killing is wrong *only* because God said: "Thou shalt not kill," God might just as easily have said: "Thou shalt kill." Would killing then have been right? To agree that it would have been right makes morality too arbitrary; but to deny that it would have been right is to assume that there are standards of right and wrong independent of God's will.²⁹

The application to a pure virtue-theorist's approach should be apparent. Julia Driver puts the issue as follows, with respect to right action:

it seems natural to ask the further question, "Why would the virtuous agent advise me to do A?" If the answer is simply that what the virtuous agent advises determines right action, independent of any other reasons or considerations, then the account seems quite capricious; if, on the other hand, there are independent reasons, then aren't those the right-making features – and then isn't what the virtuous agent advises superfluous?³⁰

Intuitively, either the actions (or approvals and disapprovals) of virtuous agents follow some set of independent standards of rightness or goodness (in which case these are fundamental, not the attitudes of virtuous agents), or else the actions and approvals of the virtuous are simply arbitrary (which, in turn, makes such an ethics arbitrary and not worthy of our concern).

What can be said in response? While there are certain dangers in appealing to the color / secondary quality model in understanding virtue theories, it will be instructive in the present case as a starting point. We are trying to find a model that allows us to hold that virtuous agents are not simply appealing to prior standards of rightness and wrongness, while also not simply arbitrarily approving of some things, disapproving of others. Consider now colors. It is considered quite plausible among many that color properties are not properly understood entirely independently of the reactions of observers. For example, if we consider objects that are yellow, it seems highly unlikely that we would group just these objects together were it not for the reactions of normal observers under normal circumstances. The same would be true with other colors. We "carve up" the world in a way that would seem rather arbitrary were it not for the existence of human observers with common, shared reactions to these objects.

Yet this is not arbitrary. Given our human visual systems, the world will affect us in common, predictable ways (and we cannot simply choose to see things otherwise). Light with a wavelength of near 660nm will appear red to us; with a wavelength of near 410nm it will appear as violet; at near 570nm it

²⁹ Peter Singer, *The Expanding Circle: Ethics and Sociobiology* (New York: Farrar, Strauss & Giroux, 1981), x.

³⁰ Julia Driver, "Virtue Theory," in James Dreier, ed., *Contemporary Debates in Moral Theory* (Malden, MA: Blackwell, 2006), 113-23; quotation from 118.

will appear to be yellow, and so on. We will carve up the world on this basis, but surely would not do so without our given perceptual faculties; for example, the range of wavelengths that is seen as red is much wider than the blue range (that is, colors do not nicely line up into equal, say, 50nm ranges. Blue is quite a narrow band of the spectrum).³¹ Outside of a narrow range of wavelengths, approximately 400-700nm, light is no longer visible to humans, and we do not divide the electromagnetic spectrum into colors outside of this range.

Here we can return to how we ought to model the virtue theorists' understanding of rightness, goodness, and other moral properties. Virtuous agents, given their psychological dispositions, will often immediately respond in certain ways to various events, actions, and so forth. As in the case of colors, we could perhaps arrive at a complex rule (or set of rules) that corresponds to their approvals and disapprovals. But crucially these rules could simply be a reflection or summary of the immediate reactions of the virtuous; the virtuous themselves need not appeal to a set of rules (or a prior goodness) in arriving at their simple reactions. For example, benevolent agents, given their psychologies, will tend to be moved by instances of pain or suffering. But there is no need to posit a prior, intrinsic badness to these pains, nor a prior set of moral rules to explain the regularity of these responses of benevolent agents.

Still, as noted above, appeal to the color analogy is not without its dangers and we now need to address a disanalogy. In the case of colors, normal observers do not typically engage in reasoning to determine what color they are seeing. But while in many cases virtuous agents will have immediate reactions (as with benevolent agents being immediately moved by suffering), they also engage in moral reasoning. They often need to weigh conflicting considerations to arrive at considered judgments – which is quite unlike what occurs in normal color perception. And even the immediate, simple reactions of the virtuous can be the result of training (as noted in section II). We arrive at a variation of the original dilemma: how do virtuous agents engage in moral reasoning? It seems that either they would appeal to the right or to the good, or their “reasoning” would be unconstrained, arbitrary, and quite possibly inconsistent.

To see a way through, we need to bear in mind that the virtuous, given their psychologies, will tend to be moved by certain states of affairs (e.g., honest agents will tend to have negative reactions to acts of lying). Similarly, they will be inclined toward embracing certain rules of thumb (given their concern with the truth, honest agents will tend to endorse rules of thumb that prohibit lying in most circumstances). But notice that these approvals and reactions are grounded in the psychologies of the virtuous, rather than being responses to explanatorily prior rightness or goodness. When there are

³¹ And there are, of course, further complications. For example, light of two different wavelengths can be perceived by the human eye as a third color.

conflicts between a virtuous agent's concerns (rules of thumb that appeal to her, particular reactions to a case at hand, etc.) she can attempt to achieve a reflective equilibrium among these various components. And notice that the reflective equilibrium she will achieve will reflect her virtuous dispositions. Intuitively, there will be a significant difference between how a Hitleresque character would arrive at a reflective equilibrium, and the equilibrium that would be achieved by Albert Schweitzer. Presumably the virtuous agent will strive for consistency (her disposition toward justice or fairness would seem to require this), thus avoiding arbitrariness or inconsistency. We thus arrive at the following: the virtuous, given their concerns and dispositions, will attempt to achieve a reflective equilibrium; the very equilibrium they arrive at will in turn be shaped by their dispositions (and again, their inclinations toward justice will ensure a desire for consistency). There is no need to appeal to an explanatorily prior rightness or goodness. We thus arrive at a way through the dilemma: the reactions of the virtuous are regular and consistent (avoiding the second horn), but are not the result of following of some prior set of rules or goodness (avoiding the first horn).

V

And this leads us to a final objection. The objector might grant that there is a concern for certain states of affairs built into the various virtues. But then we arrive at a fundamental question. How do we pick out the virtues? And here a plausible story would hold that various character traits are virtues because they include a concern for states of affairs (or action-guiding principles) that are independently valuable (or disvaluable). Benevolence is a virtue, while malevolence is not, because benevolence involves a concern for the happiness or well-being of others, where happiness is independently valuable. William Frankena provides a clear, classic presentation of the objection:

[O]ne cannot conceive of traits of character except as including dispositions and tendencies to act in certain ways in certain circumstances. Hating involves being disposed to kill or harm, being just involves tending to do just acts (acts that conform to the principle of justice) when the occasion calls. Again, it is hard to see how we could know what traits to encourage or inculcate if we did not subscribe to principles, for example, to the principle of utility, or to those of benevolence and justice.³²

The question, then, is whether we can give an account of the virtues (including a story about what makes given traits virtues) that does not boil down to having a concern for independently (and explanatorily prior) right actions, principles, or good states of affairs.

Here I would argue that such independent accounts are clearly possible, while conceding that most of the accounts that have in fact been offered would be inadequate for current purposes. Thus, appeals to developing hu-

³² William Frankena, *Ethics* 2nd ed. (Englewood Cliffs, NJ: Prentice Hall, 1973), 65.

man nature (or leading a flourishing human life) have been frequently endorsed, where the virtues are taken as essential to human nature (or leading a *eudaimon* life).³³ On such accounts we might hold that benevolence is a virtue as it is essential to developing our human nature (as social, rational animals), or because it is necessary to leading a flourishing human life. Notice that we would not necessarily need to appeal to prior morally right actions or good states of affairs in establishing the virtues. As such, we would have an independent account of the virtues that would anchor the responses to the objections discussed throughout this paper. On the other hand, most proponents of such accounts (including Hursthouse) explicitly appeal to moralized accounts of human nature; such accounts would likely not serve our present purposes of avoiding Frankena's charge.³⁴ Still, we might find a satisfactory account along these lines, as philosophers continue to develop and refine such approaches.³⁵

And, of course, other accounts are possible. One promising path would involve (roughly) appealing to those character traits that we would come to value in a society's members, given full information (about the probable effects of various traits and so on). Rather than appealing to human nature, we could appeal to those traits that we would value in people, in light of full information. Again, this would allow us to pick out traits as virtues without assuming explanatorily prior morally right actions or good states of affairs. Variations of this general approach are also available.

There is thus admittedly a rather large promissory note left to be paid in explaining what marks off traits as virtues. But this is not a circumstance unique to virtue theories,³⁶ and there are promising avenues for virtue theorists to explore in discharging such debts. Here I hope only to have cleared

³³ While there are important differences between their various accounts of human nature and flourishing, we could include as examples of this general approach: Aristotle, *Nicomachean Ethics*, Terence Irwin, trans. (Cambridge: Hackett, 1985), William D. Casebeer, *Natural Ethical Facts: Evolution, Connectionism, and Moral Cognition* (Cambridge: The MIT Press, A Bradford Book, 2003), Philippa Foot, *Natural Goodness* (New York: Oxford University Press, 2001), Rosalind Hursthouse (1999), and James D. Wallace, *Virtues and Vices* (Ithaca, NY: Cornell University Press, 1978).

³⁴ This would depend on the details of any given moralized account of human nature or flourishing – does the account make reference to explanatorily prior right actions or the good? If so, the account would not suffice for our current purposes.

³⁵ See, for example, Thomas Hurka, *Perfectionism* (New York: Oxford University Press, 1993). Hurka attempts to avoid moralizing his account of human nature, but explicitly acknowledges that the resulting perfectionism will fail to include many commonsense virtues (while suggesting that perhaps his perfectionism could be supplemented with additional moral content to capture more of commonsense morality). There may be other ways in which one might develop a non-moralized perfectionism, approaches that might capture more of commonsense morality.

³⁶ That is to say, a broad range of important theories (in ethics and elsewhere) face important challenges, and have lacunae that must eventually be filled.

the ground of several important initial objections to the claim that the virtues could be explanatorily basic in a viable moral theory.³⁷

Jason Kawall
Colgate University
Department of Philosophy
jkawall@colgate.edu

³⁷ Early versions of this paper were presented at meetings of the Pacific Division of the American Philosophical Association and of the Tennessee Philosophical Association. I would like to thank my commentators (Dan Farnham and Dina Garmong) and audiences for their comments. Thanks also to those who discussed certain ideas from this paper on the ethics blog, PEAsoup (<http://peasoup.typad.com/peasoup>) – Ben Bradley, Campbell Brown, Jamie Dreier, Josh Glasgow, Troy Jollimore, and Scott Wilson. Finally, I would like to thank Ron Sandler, Gretchen von Schwinn, and the anonymous referees of this paper for their very helpful comments.