

Moral Worth and Our Ultimate Moral Concerns

Douglas W. Portmore

Arizona State University
dwportmore@gmail.com

Abstract: Some right acts have what philosophers call *moral worth*. A right act has moral worth if and only if it manifests a set of motives/concerns that makes its agent deserving of praise for having non-accidentally done what's right. I argue that what makes an agent deserving of praise for having non-accidentally done what's right is that her act manifests an appropriate set of concerns, where the appropriateness of these concerns is a function what her ultimate moral concerns should be. Two important upshots of the resulting account of moral worth are that (1) an act can have moral worth even if it doesn't manifest a concern for doing what's right and that (2) an act can lack moral worth even if it is performed for the right reasons.

1. *Moral Worth*

Some right acts have what philosophers call *moral worth*.¹ These acts manifest the agent's virtuous motives and, thereby, do her credit.² More precisely, a right act has moral worth if and only if it manifests a set of motives/concerns that does its agent credit in that it makes her deserving of praise for having non-accidentally done what's right. Now, not all right acts have moral worth. Take, for instance, Immanuel Kant's shopkeeper (G 4:397).³ He deals honestly

¹ By 'right', I mean 'morally right'. And I use the term broadly to include both obligatory acts and supererogatory acts. Note, then, that I won't take a stand here on whether either impermissible acts or merely permissible acts (that is, permissible acts that are neither obligatory nor supererogatory) can have moral worth.

² An act manifests a set of motives/concerns if and only if we must appeal to these motives/concerns, but not to any others, to explain its performance.

³ The 'G' stands for Kant's *Groundwork*, and the citation is given by volume and page number.

with his customers, but only because it's good for business. So, although he acts rightly in dealing honestly with his customers, these acts have no moral worth in that they were motivated solely out of a selfish concern to maximize profits, a concern that would have led him to cheat his customers had they been less savvy. Thus, his acting rightly was merely fortuitous. For had it not been in his self-interest to act rightly, he would have acted wrongly.

The moral worth of an act is a function of the virtuousness of the motives/concerns it manifests as opposed to the virtuousness of the character of the agent who performs it. Thus, an act can have moral worth even if its agent has a bad character. Consider that even a stingy miser might do something generous on occasion. And if, on that occasion, what moves her is an appropriate set of concerns, her act will have moral worth despite being out of character. Likewise, someone with a good character might act from bad motives on occasion, and, when she does, her act will lack moral worth. So, whereas the moral worth of an act depends on what the agent's motives were in the given situation and whether those motives could potentially lead her to act wrongly in other situations, an agent's character has to do with whether she's disposed to have the appropriate concerns in a wide range of situations (even if not necessarily in the given situation). Thus, the stingy miser counts as having a bad character, because, in most situations, she'll care too much for her own wealth and/or too little for the welfare of others. Nevertheless, it may be that she has the appropriate set of concerns and acts generously on that rare day in which everything goes her way: she finds her morning paper on her doorstep instead of in the bushes, the barista gets her order right and even spells her name correctly, the other commuters with whom she shares the road that morning are unusually courteous, and, when she arrives at work, she learns that she's getting the promotion that she'd been hoping for. Suppose that, on that day, she buys lunch for a homeless man out of a concern that he doesn't go hungry. And this act has moral worth because it manifests an appropriate set of concerns (e.g., a greater concern for this man's welfare than for squirreling away a few extra dollars for herself), a set of concerns that made her acting rightly in this situation non-accidental. So, the fact that she has a bad character and wouldn't have acted generously or have had the appropriate concerns if this had been anything but an exceptionally good day for her doesn't detract from the moral worth of this particular act, which manifests a virtuous set of

concerns.

Now, when it comes to the moral worth of an act, it's not just the specific concerns that moved the agent to perform it that matters; her other concerns can matter as well. To illustrate, consider the following case.

The Dog-Lover: A dog-lover named Yunn protects a poodle from a boy's kick by blocking his blow with her own leg. And she does so out of a concern for the dog's welfare. Thus, she does the right thing for the right reasons. But suppose that, in this instance, Yunn had absolutely no concern for the welfare of the boy and cared only for herself and the dog.⁴ So, she would have fatally shot the boy had this been an option for her. For, in that case, she could have protected the dog without having to suffer his painful blow. But, as it was, she didn't have this option and could protect the dog only by blocking his blow with her own leg. So, her acting rightly in this instance was merely accidental. Indeed, the same set of concerns that led her to do the right thing in this situation would have led her to do the wrong thing in other relevantly similar situations, such as the one in which she had the option of shooting the boy.⁵

The lesson, I take it, is that we must look not only at the specific motive/concern that led her to act as she did, but also at her other pertinent concerns. And I'll be assuming that a lack of

⁴ I'm claiming only that these were her concerns on this particular occasion. Thus, I'm not making any claim about her character or what sorts of relatively stable sets of concerns she would have across various counterfactual situations.

⁵ Julia Markovits (2010, 210) talks about a somewhat similar case. In her case, a fanatical dog-lover saves several strangers at great risk to herself. But, given her fanatical love for dogs, she would not have saved them had the choice been between saving them and saving her dog. Markovits claims that, assuming that the dog-lover's preference for saving her dog over the strangers is the result of her having too much concern for her dog rather than too little concern for the strangers, her act of saving the strangers has moral worth despite the fact that her excessive concern for her dog would have led her to do the wrong thing in other situations. Now, I agree with Markovits on this point provided that we're assuming that what interests us is whether this woman is willing to sacrifice her own interests for the sake of promoting the much greater interests of others. But, as I'll argue below, there may be other contexts in which what interests us is whether the woman's concern for her dog is excessive. And, in those contexts, her act would not count as having moral worth.

concern for something can count as a concern; it just counts as a “zero concern” for that thing. Thus, in *The Dog-Lover*, we must consider not only Yunn’s concern for both herself and the dog, but also her zero concern for the boy. For she would have been led by this set of concerns to act wrongly in a situation in which she had the option of shooting the boy. Thus, it was lucky that she didn’t have this option. And, so, we should think that an act’s moral worth depends not merely on the agent’s motivating reason for performing it, but on all her pertinent concerns.

Now, for any agent in a given situation, her *pertinent* concerns are all and only those that will (or would) determine whether she acts rightly in this and other relevantly similar situations. Thus, in *The Dog-Lover*, Yunn’s zero concern for the boy’s welfare is pertinent given that it would have led her to act wrongly in the relevantly similar situation in which she had the option of shooting the boy and everything else was the same. In that situation, she would have acted wrongly. Thus, her acting rightly in the given situation was merely fortuitous. Likewise, if Yunn had had a concern for the boy but not for dogs with spots, her acting rightly would have counted as merely accidental. For a zero concern for dogs with spots would have led her to refrain from acting rightly (that is, to refrain from blocking the boy’s blow) in the relevantly similar situation in which the boy was about to kick a Dalmatian rather than a poodle.

Of course, not every concern that determines whether an agent would act rightly in some other situation counts as pertinent. For, when it comes to pertinence, it’s only the *relevantly similar* situations that matter. Thus, even if Yunn had had a concern to prevent Muslims from immigrating to the U.S., this wouldn’t itself prevent her act of protecting the dog from having moral worth. Although this concern would have led her to do the wrong thing in a situation in which she had the opportunity to prevent a Muslim with a compelling asylum claim from immigrating to the U.S., this situation isn’t relevantly similar to the one at hand: one in which she has to choose whether to protect a dog from harm. That said, there’s no simple and straightforward way of spelling out what the relevantly similar situations are. We may wonder, then, whether a situation in which a boy is about to beat a snake with a stick is relevantly similar to the one described in *The Dog-Lover*. And, so, we may wonder: Would Yunn’s having a concern for the welfare of both dogs and humans but not for the welfare of snakes make her

acting rightly in *The Dog-Lover* count as merely accidental? The answers to such questions depend on the context in which they're being asked and on what's taken to be the relevant similarity relation in that context. Thus, we could imagine one context in which we're concerned with whether Yunn is speciesist and, thus, whether she has a concern for the welfare of all sentient beings and not just for the welfare of her fellow mammals. And, in that case, her acting rightly in this situation would count as merely accidental. But we could also imagine a context in which we're merely concerned with whether Yunn is sufficiently altruistic and so willing to sacrifice her own welfare for sake of promoting the welfare of some other being. And, in that case, her acting rightly in this situation would count as non-accidental.

As I see it, this sort of contextualism is a feature rather than a bug. For one, moral worth has to do with whether or not an agent was just lucky to have acted rightly, and, in general, whether someone counts as lucky is contextually determined. Take, for instance, the person who becomes infected with the Ebola virus and survives. On the one hand, we might consider her to be rather lucky, as only one in three survive such an infection. On the other hand, we might consider her to be extremely unlucky, as Ebola infections are exceeding rare. For another, this sort of contextualism helps us to explain why there tends to be so much intractable disagreement about various cases in the literature about moral worth. The disagreement is intractable, because there's no set answer in the abstract as to what the relevant similarity relation is. It just depends on the context and what's assumed to be relevant in that context. I'll have more to say about this below. But, for now, I should just admit that I have nothing interesting to add to the existing literature on how to identify the relevant similarity relation in a given context.⁶ So, in the rest of the paper, I will just focus on cases in which I suspect there will be wide agreement as to what the relevant similarity relation is.

As we've just seen, a central feature of an act with moral worth is that it is praiseworthy in virtue of its manifesting a set of pertinent motives/concerns that makes the agent's having

⁶ But see the existing literature concerning similarity relations and each of the following: the counterfactual account of harm (e.g., Bradley 2009), the counterfactual account of causation (e.g., Schaffer 2000), and the truth conditions for counterfactuals (e.g., Lewis 1973).

acted rightly non-accidental.⁷ And, as we learned from the shopkeeper case, an act counts as merely accidentally right if it manifests a selfish motive given that such a motive would lead one to act wrongly in relevantly similar situations in which it's not in one's self-interest to act rightly.⁸ And, as we learned from *The Dog-Lover*, an act can count as merely accidentally right even if it stemmed from a good motive (e.g., a concern for the welfare of a dog) if that motive was, nevertheless, part of a set of pertinent concerns that would have led one to act wrongly in other relevantly similar situations. But we also learned both that it's only the pertinent concerns that matter and that not all concerns are pertinent. Thus, Yunn's act wouldn't lack moral worth simply because she had a concern that Muslims not immigrate to the U.S., for this concern isn't a pertinent one—at least, not in any obvious sort of context. Also, it's important to realize that even the best motives can lead one to act wrongly when one is misinformed. So, what we're really interested in is whether the agent's pertinent concerns would ever lead her to act wrongly in a situation in which she is relevantly informed. And, so, we should adopt the following criterion for an act's having moral worth.

The Non-Accidentality Criterion: A right act has moral worth if and only if it manifests a set of pertinent concerns that would never lead its agent to act impermissibly in a situation in which she's relevantly informed.⁹

⁷ There seems to be consensus on this point. See, for instance, Arpaly (2002, p. 225), Herman (1981, p. 366), Isserow (Forthcoming), Johnson King (2020, p. 192), Markovits (2010, pp. 206 & 211), Singh (Forthcoming), and Sliwa (2016, pp. 395–398).

⁸ I'm assuming that, in every moral context, we'll be concerned with whether an agent is fulfilling her moral duty only because doing so happens to align with her self-interest. And this, I think, is why there seems to be no disagreement concerning whether, in Kant's example, the shopkeeper's act of treating his customers honestly counts as merely accidentally right given that it was motivated solely out a concern for his own self-interest.

⁹ In many instances, an agent will be led to acquire new knowledge out of a concern, say, for people's safety. For instance, before handing a prop gun to an actor, a stage hand with such a concern would be led to determine whether it's loaded with blanks or live ammunition. Of course, if the stage hand were fully informed, there would be no need for her to check the gun so as to acquire the knowledge that she already has. But if she's relevantly (though not fully) informed, she would be led by such a concern to check the gun. So, given that being fully informed can affect what an agent ought to do even in the fact-relative sense of 'ought', I'll be concerned with agents who are *relevantly* informed as opposed to *fully* informed.

With this criterion in hand, I plan, first, to assess various proposed accounts of moral worth and, second, to propose a new account. I'll start by looking at two simple views. Although no contemporary philosopher accepts such a simple view, most contemporary views can be seen as more sophisticated versions of these two. Nonetheless, I believe that both views need more than mere revision, as I believe that both are fundamentally flawed. I'll explain at the end of the next section what this fundamental flaw is. And this will lead me to introduce a new concept—the concept of an ultimate moral concern—in section 3. I will then, in section 4, employ this concept in developing a new account of moral worth and show how this account compares favorably to its rivals. Lastly, in section 5, I conclude by explaining the extent to which this new account is and isn't trivial.

2. Two Opposing Accounts of Moral Worth: The Kantian View and the Humean View

Most contemporary accounts of moral worth stem from the views of either David Hume or Immanuel Kant. I'll present only the simplest version of each, and I make no claim as to their historical accuracy. My aim is merely to lay out the two most basic points of view from which most contemporary views have been developed. I'll start with the view that's been inspired by Kant and his thought that moral worth attaches to right actions that are performed simply because they are right—i.e., actions motivated “from duty” or “from respect for the law” (G 4:396–401). This has inspired the following account of moral worth.

The Kantian View: A right act has moral worth if and only if it manifests a non-instrumental concern for performing acts that are right. In other words, assuming the Humean theory of motivation, a right act has moral worth if and only if its agent was motivated to perform it out of both a non-instrumental desire to do what's right and the belief that were she to perform this act she would do what's right.

Note that, on the Humean theory of motivation, an “agent A at t has a motivating reason to φ only if there is some ψ such that, at t , A desires to ψ and believes that were he to φ he would ψ ” (Smith 1987, p. 37). I appeal to this theory in my formulations of these two simple accounts of moral worth, not because I think that it’s correct, but because its simplicity allows me to more perspicuously illustrate how these two views differ from each other.

Now, there are some clear counterexamples to the Kantian View. Here’s one.

The Empathic: A man named Christoforos believes that chimpanzees are not morally considerable beings, for he thinks that they can be neither harmed nor wronged. Yet, he finds himself empathizing with the apparent plight of a chimpanzee that has just been captured by poachers. Intellectually speaking, he doesn’t believe that there are any genuine feelings underlying the chimpanzee’s outward “signs” of distress. But, on an emotional level, he accurately perceives that the chimpanzee is in genuine distress. And, given these perceptions, he empathizes with the chimpanzee’s plight having once been held captive himself.¹⁰ So, when the opportunity arises, he’s moved to help the chimpanzee escape back into the wild out of a concern to alleviate the chimpanzee’s apparent distress. And this is his sole motive, for he doesn’t think that his helping the chimpanzee escape is the right thing to do. (Nor does he think that it’s the wrong thing to do.) Afterwards, he’s tempted to just walk away from the situation. Yet, he ends up reporting the poachers to the authorities out of both a concern for doing what’s right and the belief that reporting lawbreakers to the authorities is one’s duty. What’s more,

¹⁰ I believe that, through our emotional experiences, we can apprehend important truths. And these experiences provide us with evidence for these truths. What’s more, they can represent the world as being one way even while our avowed beliefs represent the world as being another way. And sometimes it’s our emotions rather than our avowed beliefs that accurately represent the way the world is. (See Furtak 2018, especially chap. 3.) That’s what I take to be going on with Christoforos. Through his empathic response to the apparent signs of the chimpanzee’s distress, he accurately represents the world as being one in which the chimpanzee is suffering, and yet, through his beliefs (or, at least, in terms of what he’s willing to assent to), he inaccurately represents the world as being one in which the chimpanzee is not suffering. Fortunately, his actions seem to be guided by what his emotions are telling him rather than by what his avowed beliefs are telling him.

he facilitates the authorities taking the poachers into custody safely by hiding their guns, and he does so out of a concern for the welfare of both the poachers and the authorities and despite knowing that this entails some added risk to himself.¹¹

On the Kantian View, Christoforos's act of helping the chimpanzee escape back into the wild doesn't have moral worth because it fails to manifest a non-instrumental concern for doing what's right. Intuitively, though, it seems to have moral worth, for it seems to manifest an appropriate set of concerns—one that made his acting rightly in this situation non-accidental. After all, he had all the pertinent concerns: including a concern not only to alleviate the chimpanzee's apparent distress, but also to safeguard the welfare of both the poachers and the authorities. And he even had a concern for doing what's right. What's more, the magnitude of each of his concerns was, we'll stipulate, at the appropriate level. And, given all this, he would never be led to act wrongly by such a set of concerns in any situation in which he was relevantly informed. So, we should think that, contrary to what the Kantian View implies, his act of helping the chimpanzee has moral worth. Thus, the Kantian View violates the accidentality criterion and should be rejected.

The other leading inspiration for accounts of moral worth is David Hume. According to Hume, "no action can be virtuous, or morally good, unless there be in human nature some motive to produce it, distinct from a sense of its morality" (T 3.2.1.7).¹² This suggests the following account of moral worth.

¹¹ More commonly, philosophers (such as Arpaly 2002) cite the case of Huck Finn as a putative counterexample to the Kantian View. I prefer this example, because it's unclear whether Huck's concerns are appropriate. For one might argue that Huck should have a greater concern for doing what's right and/or a lesser concern for being loyal to a friend. And if Huck does, say, have too great a concern for being loyal to a friend, then he might be led by such a concern to act wrongly in other situations in which he's relevantly informed. For instance, as Zoë Johnson King (2020, p. 195) worries, Huck might be led by such a concern to help a friend elude the authorities even when that friend is a serial killer on the run rather than a fugitive slave.

¹² The 'T' refers to Hume's *A Treatise of Human Nature* and the citation is given by book, part, section, and paragraph number.

The Humean View: A right act has moral worth if and only if it manifests a non-instrumental concern for performing acts that have some right-making feature RMF. In other words, assuming the Humean theory of motivation, a right act has moral worth if and only if the agent was motivated to perform it out of both a non-instrumental desire to perform an act that has RMF and the belief that were she to perform this act she would perform an act that has RMF.

This view is also subject to counterexample, though what sort of counterexample it is subject to depends on which of the following two versions of it we accept.

On the fundamentalist version, RMF must refer to the act's *fundamental* right-making feature. Thus, if maximizing act-utilitarianism is correct, 'some right-making feature RMF' must refer to 'the feature of maximizing aggregate utility'.¹³ Given this view and the Humean theory of motivation, an act will have moral worth if and only if the agent is motivated to perform it out of both a non-instrumental desire to maximize aggregate utility and the belief that were she to perform this act she would maximize aggregate utility.

On the non-fundamentalist version, by contrast, RMF can refer to any right-making feature, whether it's fundamental or not. To illustrate, assume that maximizing act-utilitarianism is correct and suppose that I would maximize aggregate utility if and only if I were to push the button that's in front of me, for pushing this button is my only option for

¹³ Nomy Arpaly and Timothy Schroeder seem to endorse the fundamentalist version of the Humean View. They hold that acting for the right reasons (which they equate with acting in a way that has moral worth) is just a matter of having intrinsic desires that are instances of good will. And they say that "for an intrinsic desire to be an instance of good or ill will the content of the desire must be something one has a *pro tanto* moral reason to do or avoid and this content must be presented by concepts that would allow the individual in question to trivially deduce that it is necessarily an instance of MAXIMIZING HAPPINESS, or RESPECTING PERSONS, or whatever the correct normative theory distinguishes as the right or good as a whole." (2014, p. 167). Thus, as Arpaly explains in another work, "a morally worthy action stems from a commitment to the right and the good *correctly conceptualized*. If utilitarianism has the right account of the features that make actions right then the agent performing a morally worthy action conceives of her action as maximizing utility, and is committed to maximizing utility so conceived" (2015, p. 87).

saving many lives. Also, for the sake of argument, assume that the Humean theory of motivation is correct. Given these assumptions, the non-fundamentalist version of the Humean View implies that my act of pushing the button will have moral worth if I am motivated to perform it out of both a non-instrumental desire to maximize aggregate utility and the belief that were I to perform this act I would maximize aggregate utility. But it also implies that my act of pushing the button will have moral worth if I am instead motivated to perform it out of both a non-instrumental desire to save many lives and the belief that were I to perform this act I would save many lives. For, in this instance, saving many lives is what would maximize aggregate utility. Thus, saving many lives is what makes my pushing the button the right thing to do. It's just that this is derivatively so. Thus, on the non-fundamentalist version of the Humean view, 'some right-making feature RMF' can refer to 'the feature of maximizing aggregate utility', 'the feature of saving many lives', or any other right-making feature.

Unfortunately, both versions of the Humean View are problematic. The problem with the non-fundamentalist version is that it gets the wrong result in cases like *The Dog-Lover*. For this version of the Humean View holds that Yunn's act has moral worth given both that she had a non-instrumental concern to prevent the dog from getting hurt and that this is what makes her blocking the boy's kick with her own leg the right thing to do. Of course, it's not what fundamentally makes it right. For assuming (merely for the sake of illustration) that maximizing act-utilitarianism is correct, what fundamentally makes it right to block the boy's blow is that doing so would maximize aggregate utility.¹⁴ Nevertheless, that the act prevents the dog from getting hurt is what derivatively makes it right given that preventing the dog from getting hurt is what would maximize aggregate utility in the given situation. So, on the non-fundamentalist version, Yunn's act has moral worth. But, as we saw above, Yunn's act is merely accidentally right given that she has zero concern for the boy's welfare. And, thus, she would have shot the boy were that an option. Thus, the non-fundamentalist version of the Humean View violates the non-accidentality criterion and should be rejected.

¹⁴ I'm assuming that the boy's blow would cause a lot more harm to the dog if not blocked by Yunn than it would cause to Yunn if blocked by her.

The reason the non-fundamentalist version of the Humean View violates the non-accidentality criterion is that it fails to capture the counterfactual reliability that's required for moral worth. Consequently, it allows that the set of concerns that you'll be praising an agent for are ones that would lead her to act wrongly in other situations in which she is relevantly informed. This is because the only way to ensure counterfactual reliability is to look not only at the agent's motivating reason and whether it was good, but also at whether the agent had all the other pertinent concerns and in the correct proportions. This is because whether an act is permissible depends not merely on whether it has some good feature (which might be the basis for an agent's motivating reason for performing it), but also on whether it has any outweighing bad feature. Thus, doing something to protect a dog from a boy's kick is permissible when it involves blocking that kick with one's own leg, but not when it involves shooting him before his kick has a chance to connect. So, to ensure counterfactual reliability while adopting the Humean View, we must adopt the fundamentalist version of the Humean View. After all, it's only a non-instrumental concern to perform acts that have the *fundamental* right-maker that will ensure that one never does wrong when relevantly informed.

But the fundamentalist version of the Humean View is also unacceptable. The problem is that it makes moral worth too hard to come by. As Daniel Star has pointed out, people rarely conceptualize their actions as meeting some fundamental moral criterion.¹⁵ Even self-professed maximizing act-utilitarians rarely conceptualize what they're doing as maximizing aggregate utility. And this is a good thing. Because there is, as Bernard Williams (1981) has pointed out, something very wrong with someone who is, say, moved to kiss her partner out of a concern to maximize aggregate utility rather than simply as a result of her affection for her. What's more, it seems that an act can have moral worth even if its agent wasn't motivated by a non-instrumental desire to do that which maximizes aggregate happiness and the belief that by acting this way she will maximize aggregate happiness (or, substitute here whatever the fundamental right-maker is). For it seems sufficient that she was motivated out an appropriate level of concern for the welfare of each of the sentient beings involved.

¹⁵ See Star (Forthcoming). See also Howard (Forthcoming).

In any case, *The Empathic* seems to be a clear counterexample to the fundamentalist version of the Humean View. Christoforos wasn't motivated out a non-instrumental concern for anything such as maximizing aggregate utility, abiding by the ideal code of rules, or acting in accordance with the categorical imperative. Rather, he was motivated simply out of a desire to alleviate the chimpanzee's apparent distress. What's more, he had all the other pertinent concerns. For he cared about the welfare of both the poachers and the authorities. And he even cared about doing what's right. He just didn't have an additional concern for doing what would maximize aggregate utility, conceived as such. But caring about each individual and in the correct proportions (in, say, proportion to the amount of welfare that's at stake for each of them) will unerringly lead him to maximize aggregate utility in any situation in which he is relevantly informed. Thus, we should reject the fundamentalist version of the Humean View, for it too violates the non-accidentality criterion.

So, we should reject both the Kantian View and the Humean View. We should reject the Kantian View because it gets the wrong answer in *The Empathic*. And we should reject the Humean View because it gets the wrong answer in either *The Empathic* or *The Dog-Lover*, depending on whether we're considering its fundamentalist or non-fundamentalist version.¹⁶ Now, there have been several attempts to salvage some version of these two views. But I doubt that either can be salvaged, for they both go wrong in a very fundamental way. Specifically, they both go wrong in failing to acknowledge that all and only those right acts that manifest an appropriate set of concerns have moral worth.

More specifically, the Kantian goes wrong in insisting that acts with moral worth must manifest a non-instrumental concern for doing what's right, when, arguably, having a *non-instrumental* concern for doing what's right is inappropriate. For it seems that we shouldn't care

¹⁶ Jessica Isserow (Forthcoming) has recently defended a pluralist proposal according to which "it is necessary and sufficient for an agent's action to have moral worth that she be motivated either by the consideration that her action is morally right, or by the considerations that explain why her action is morally right." I believe that her view will suffer the same fate as the Humean View, because either "the considerations that explain why the agent's act is right" will refer exclusively to the act's fundamental right-maker or it won't. If it does, then it will get the wrong result in *The Empathic*. And if it doesn't, then it will get the wrong result in *The Dog-Lover*.

about doing what's right for its own sake. Rather, we should care about doing what's right only as a means to doing right by whatever it is that we should ultimately care about. Nathan Howard (Forthcoming) makes the point nicely.

Acting from a desire for rightness as such...is a little like desiring to get the cheap plastic trophy without caring about whether you're the champion. The trophy is worth getting only because it represents the verdict that you're the champ. Therefore, desiring the cheap plastic trophy as such fetishizes the trophy; it displaces your desire from its fitting object, namely, the end of being the champ. Likewise for rightness. If rightness is worth caring about, it is only derivatively so, in virtue of its connection to the ends at which morality properly aims like equality, welfare, and the care that we owe to our friends, family, and fellow humans.

And the Humean View goes wrong in denying that an act can have moral worth in virtue of manifesting an *instrumental* concern for doing what's right, when, arguably, such a concern is entirely appropriate. For sometimes we either don't know what we should ultimately care about or don't know how to do right by that which we should ultimately care about. But we may, nevertheless, know what the right thing to do is—as a result of, say, the reliable testimony of someone who we can trust to know what's right in such situations. And, in such instances, we should care about doing what's right as a means to doing right by whatever it is that we should ultimately care about. Indeed, as Paulina Sliwa points out, “conative states with moral content (e.g., a desire to do what's right) are essential for doing the right thing in the face of moral uncertainty” (2016, p. 408). What's more, there will be times when we are tempted to do wrong because we don't directly care enough about what we should. And, in such cases, we can care about doing what's right as a proxy for directly caring about what we should. Yet, according to the Humean View, an act has moral worth only if it manifests a non-instrumental concern for its right-making features, and an instrumental concern for an act's being right isn't the same as having a non-instrumental concern for its right-making features.

So, as we've just seen, both the Kantian View and the Humean View fail to accommodate the plausible idea that all and only those right acts that manifest an appropriate set of concerns have moral worth, where the appropriateness of a concern—both in terms of its magnitude and in terms of its being either instrumental or non-instrumental—is determined by

what our ultimate moral concerns should be. The Kantian View fails in requiring us to have an inappropriate concern (that is, a non-instrumental concern for doing what's right), and the Humean View fails in prohibiting us from having an appropriate concern (that is, an instrumental concern for doing what's right). Of course, in making these arguments, I've relied heavily on the notion of what we should ultimately be concerned about. Unfortunately, this notion has been undertheorized. So, at this point, I will briefly digress from our discussion of moral worth to further explicate the notion. I'll then, in the next section, show how we can exploit this notion to develop an account of moral worth that avoids the problems associated with both the Kantian View and the Humean View.

3. *Our Ultimate Moral Concerns*

A complete moral theory should provide us not only with accounts of both what's right and what's good, but also with an account what our ultimate moral concerns should be. Such an account would tell us what we morally ought ultimately to care about and, thus, what we morally ought to be ultimately aiming to achieve or, at least, what we morally ought to be hoping is achieved—and, if not by us, then by others. For one, the question of what we morally ought ultimately to care about is itself an important moral question and, thus, one that we should expect a complete moral theory to answer for us. For another, whether the ultimate moral concerns that a moral theory prescribes for us is consistent with our being motivated to do as its criterion of rightness directs us to act will determine whether that theory is *incoherent*—incoherent in that its criterion of rightness sometimes permits (or, even worse, requires) agents to act in ways that they know won't optimally advance the ultimate moral concerns that the theory prescribes to them. And this is important, because we should reject any theory that's incoherent.

To illustrate, consider the incoherence objection to rule-consequentialism. According to rule-consequentialism's criterion of rightness, an act is morally permissible if and only if it accords with the ideal code of rules. Now, some have worried that the ideal code will be extensionally equivalent to act-consequentialism and that, therefore, rule-consequentialism will

collapse into act-consequentialism. But, as Brad Hooker (2000) has pointed out, rule-consequentialism can avoid collapsing into act-consequentialism provided it holds that the ideal code is the one whose internalization by the overwhelming majority of everyone everywhere in each new generation would maximize the good. But, in avoiding the collapse worry, a new worry arises. For if the ideal code isn't extensionally equivalent to act-consequentialism, then there will be instances in which rule-consequentialism permits (or even requires) an agent to abide by the ideal code although she knows that her doing so won't maximize the good. And, so, if a complete version of rule-consequentialism holds both that agents must adopt maximizing the good as their ultimate moral concern and that agents are sometimes permitted (or even required) to abide by the ideal code even when they know that doing so won't maximize the good, then it will be an incoherent theory in that it will sometimes require agents to act in ways that they know won't optimally advance the ultimate moral concerns that the theory gives them. In other words, it will be incoherent in that it will require agents to have an incoherent set of motives. On the one hand, they will be required to internalize a set of motivations that will lead them to abide by the ideal code even when they know that doing so won't maximize the good. And, on the other hand, they will be required to adopt maximizing the good as their *ultimate* moral concern, such that they will be concerned with abiding by the ideal code only as a means to maximizing the good. But if they're concerned with abiding by the ideal code only as a means to maximizing the good, then they won't be motivated, as required, to abide by the ideal code even when they know that doing so won't maximize the good.

Hooker's response to the objection is to deny that a complete version of rule-consequentialism must give each agent the ultimate moral concern of maximizing the good. He says, "rule-consequentialists need not have maximizing the good as their ultimate moral goal" (2000, p. 101). He holds that rule-consequentialism is itself committed only to both a certain conception of the good (which tells us how to assess the goodness of various codes of rules) and a certain conception of the right (which tells us how to assess the rightness of acts in terms of the goodness of the various codes of rules that either permit or prohibit them), but not to any particular conception of what our ultimate moral concerns should be. So, Hooker, qua rule-consequentialist, can deny that agents should have maximizing the good as their ultimate moral

concern and hold instead that they should have ensuring that their acts are impartially defensible as their ultimate moral concern.¹⁷ And that's exactly what he does. What's more, this allows him to avoid the incoherence objection, because there is nothing incoherent about a theory that holds both that agents should have acting only in ways that are impartially defensible as their ultimate moral concern and that agents should abide by the ideal code even when they know that doing so won't maximize the good. For such a theorist can just claim that acting in accord with the ideal code ensures that one's acts are impartially defensible even when those acts fail to maximize the good.

So, it's crucial that a complete moral theory tell us what our ultimate moral concerns should be so that we can then determine whether it's incoherent. Now, admittedly, most moral philosophers have ignored this aspect of moral theorizing. They have tended to be content merely to specify each theory's criterion of rightness. Or if they go beyond that, they do so only to include a theory of the good, a theory of virtue, and/or a decision procedure. They (with, perhaps, Hooker being the singular exception) don't specify what agents' ultimate moral concerns should be. But this is clearly a mistake given that it's not just rule-consequentialism that's potentially subject to an incoherence objection. All theories are potentially subject to such an objection. Take, for instance, maximizing act-utilitarianism. It will be incoherent if it holds that agents should have an ultimate moral concern for ensuring that each sentient creature has as much utility as possible. To see why, note both that (1) maximizing act-utilitarianism holds that an act is permissible if and only if there is no alternative act that would produce a greater sum of utility than it would and that (2) although some infinities are, in some sense, larger than others, the sum of a denumerably infinite number of locations with 2 hedons each is not greater than the sum of an equal number of locations with 1 hedon each (see Kagan & Vallentyne 1997). Given these assumptions, maximizing act-utilitarianism implies that you would be permitted to φ and thereby provide an infinite number of sentient creatures with 1 hedon each even if you could instead have ψ -ed and thereby provided each of these creatures with instead 2 hedons

¹⁷ Perhaps, even this isn't what our ultimate moral concern should be. For perhaps we should be concerned with our acts being impartially defensible only because we should be concerned to show respect for people's humanity, which, perhaps, requires ensuring that we act only in ways that are impartially defensible.

each. Such a theory would be incoherent, because it permits you to act in a way that you know won't optimally advance the ultimate moral concern that it gives them. For it permits you to ϕ even though you know that ϕ -ing won't optimally advance the ultimate moral concern that you have for ensuring that each sentient creature has as much utility as possible.¹⁸

Of course, some may doubt whether I've actually identified a new and important notion. For some may suspect that this notion is just equivalent to some more commonly discussed notion. For instance, some may suspect that what I'm calling "our ultimate moral concerns" is just equivalent to what Derek Parfit calls "our theory-given aims." But, on Parfit's view, our theory-given aims are provided by the theory's criterion of rightness. Specifically, Parfit holds that our theory-given aims are just that our acts have those features that the criterion identifies as right-making and lacks those features that it identifies as wrong-making. Thus, Parfit says, "suppose that, on some theory, five kinds of act are totally forbidden. This theory gives to each of us the aim that he never acts in these five ways" (1984, p. 3). Likewise, he claims that act-consequentialism gives each of us the aim "that outcomes be as good as possible" (1984, p. 24). But our ultimate moral concerns needn't be dictated by a theory's criterion of rightness. For instance, on Hooker's version of rule-consequentialism, agents should *not* have an ultimate moral concern for performing acts that are in accord with the ideal code. Rather, they should have an ultimate moral concern for performing acts that are impartially defensible and, thus, care about abiding by the ideal code only as a means to ensuring that their acts are impartially defensible. So, what I'm calling "our ultimate moral concerns" is not what Parfit calls "our theory-given aims."

The notion of an ultimate moral concern is also distinct from the notion of rightness. What it means for an act to be permissible is that an agent would not be blameworthy for responsibly performing it. But this needn't be what an agent should ultimately care about. As

¹⁸ Interestingly, the maximizing act-utilitarian could borrow a play from Hooker's playbook and claim that, qua maximizing act-utilitarians, they needn't hold that agents should have an ultimate moral concern for ensuring that each sentient creature has as much utility as possible. They could, as Hooker does, hold that agents should instead have as their only ultimate moral concern that they ensure that their acts are impartially defensible and then hold that an act will be impartially defensible so long as there is no alternative that would produce a greater sum of utility.

we've seen, a complete moral theory, such as Hooker's rule-consequentialism, may hold that what an agent should ultimately care about is whether her acts are impartially defensible rather than whether she would be blameworthy for responsibly performing them. Indeed, the latter seems overly self-centered as far as an ultimate moral concern goes.

Lastly, the notion of an ultimate moral concern is distinct from the notion of goodness. What it means for something to be good is for it to be either better than some contextually supplied benchmark or better than a sufficient number of others in some contextually supplied comparison class (Schroeder 2008). And what it means for one thing to be better than another is for it to be fitting for an impartial spectator to prefer it to the other (Suikkanen 2008). But it may be that an agent ought to care more about one thing than another even though it would be fitting for an impartial spectator to prefer that other to it. It may be, for instance, that what I should ultimately care about is that *my* acts are impartially defensible such that I should care more about refraining from performing even one indefensible act than about preventing you from performing two indefensible acts. Still, it could be fitting for an impartial spectator to prefer that I perform one indefensible act so as to prevent you from performing two indefensible acts. So, one's ultimate moral concerns needn't be dictated by what's good, better, or best.

We've seen, then, that the notion of an ultimate moral concern is distinct from various more commonly discussed notions such as the notions of what's right and what's good. Yet, as I've shown, the notion is crucial to moral theorizing, for we should reject theories that give an account of what our ultimate moral concerns should be that makes it incoherent. In the next section, I return to the issue of moral worth and show how this notion can also help us to develop a plausible alternative to both the Kantian View and the Humean View.

4. The Concerns View and How It Compares to Its Rivals

I can now state what I take to be the correct account of moral worth.

The Concerns View: A right act has moral worth if and only if it manifests an

appropriate set of concerns—that is, a set that includes all pertinent concerns, each of which must be both qualitatively and quantitatively appropriate, which in turn is a function of what the agent’s ultimate moral concerns should be.

Remember that the pertinent concerns are all and only those that will (or would) determine whether the agent acts rightly in this and other relevantly similar situations—the relevantly similar situations being determined by the context. And whether a given concern is both qualitatively and quantitatively appropriate is a function of what the agent’s ultimate moral concerns should be. To illustrate, suppose that acting rightly is not something that an agent should have an ultimate moral concern for. In that case, she should be concerned with acting rightly only insofar as acting rightly is a means to furthering that which she should have an ultimate moral concern for. And, thus, her concern for acting rightly will be qualitatively appropriate only if it’s instrumental as opposed to non-instrumental. What’s more, it will be quantitatively appropriate only if its magnitude is proportionate to the extent to which her acting rightly is, in the given situation, a means to her furthering that which she should have an ultimate moral concern for.

To take another example, imagine that an agent should have an ultimate moral concern for promoting each existing individual’s utility (i.e., welfare) but not for promoting the overall sum of the utility in the universe. In that case, it would be inappropriate for her to care about promoting the overall sum except as a means to promoting the utility of existing individuals. Thus, it would be inappropriate for her to want to bring happy individuals into existence for the sake of increasing the overall sum of utility, but appropriate for her to want to do so as a means to promoting the utility of existing individuals, as where they would derive utility from our bringing these happy individuals into existence. Let’s further assume that this agent should care just as much about n hedons of utility for one individual as she does about n hedons of utility for any other individual. Thus, how intensely she should want to promote the utility of an existing individual by n hedons should depend on how great the number n is and not who that individual is.

Now that we have a sense of how an agent's ultimate moral concerns determine the appropriateness of the set of concerns that her act manifests, we can look at how the Concerns View deals with various cases, starting with *The Dog-Lover*. On the Concerns View, Yunn's act lacks moral worth given that she lacks a pertinent concern: a concern for the welfare of the boy. For it's plausible to suppose that, for each sentient being, Yunn should have an ultimate concern for promoting that being's welfare. Yet, she has zero concern for the boy's welfare. What's more, this concern is a pertinent one given that it would lead her to act wrongly in other the relevantly similar situations, such as the one in she has the option of shooting the boy. And since Yunn's act of blocking the boy's blow fails to manifest an appropriate set of concerns, the Concerns View rightly implies that it lacks moral worth. And there's a lesson to be learned here.

Lesson 1: A right act can lack moral worth even if it was performed for the right reasons.

After all, Yunn did perform the right act, and she did so for the right reasons (i.e., to protect the dog). Nevertheless, her act lacks moral worth, for it was merely fortuitous that she did the right thing. Had the situation been slightly different, she would have been led by the same set of concerns to have acted wrongly. And, thus, we should reject those rivals to the Concerns View that insist that moral worth is a matter of acting for the right reasons—views such as those defended by Howard (Forthcoming) and Markovits (2010, p. 205).

Let's turn now to *The Empathic*. On the Concerns View, it's not only Christoforos's act of notifying the authorities, but also his act of helping the chimpanzee escape that has moral worth. Both have moral worth because both manifests an appropriate set of concerns. In notifying the authorities and hiding the guns, Christoforos manifests a concern (an instrumental concern) for doing what's right as well as a concern (a non-instrumental concern) for promoting the welfare of each of the sentient beings involved. And these concerns are all quantitatively appropriate—or, so, I've stipulated. Likewise, his act of helping the chimpanzee escape manifests an appropriate set of concerns. Of course, it doesn't manifest a concern for doing what's right given that he doesn't think helping the chimpanzee escape is right; he thinks,

rather, that it's morally neutral. Nevertheless, his act does manifest an empathetic concern for alleviating the chimpanzee's apparent distress. And this concern is entirely appropriate given that he should have an ultimate moral concern for the chimpanzee's welfare and, thus, derivatively, for doing what would alleviate its apparent distress. So, in this case too, the Concerns View gets the intuitive verdict. What's more, the Concerns Views provides us with another important lesson.

Lesson 2: An act can have moral worth even if it doesn't manifest a concern for doing what's right.

Thus, we should reject those rivals to the Concerns View that insist that, for an act to have moral worth, it must manifest a concern for doing what's right—views such as those defended by Herman (1981), Johnson King (2020), Singh (Forthcoming), and Sliwa (2016).

Unlike the Kantian View, the Concerns View allows that an act can have moral worth even if it doesn't manifest a concern (instrumental or non-instrumental) for doing what's right. This is advantageous for two reasons. First, there are many instances in which one should not be moved by a concern for doing what's right given that this would involve one thought too many. For instance, it would, as Michael Stocker has pointed out (1976, p. 463), be disturbing to learn that the agent was motivated to rescue a child out of a non-instrumental concern to do her duty rather than out of a concern for the child (see also Williams 1981, p. 18). Second, an agent whose sole or primary motive for rescuing a child is a non-instrumental concern for doing her duty seems to have a moral fetish. For, as Michael Smith has noted, "good people care non-derivatively about honesty, the weal and woe of their children and friends, the well-being of their fellows, people getting what they deserve, justice, equality, and the like, not just one thing: doing what they believe to be right, where this is read *de dicto* and not *de re*" (1994, p. 75).

The Concerns View not only allows that an act can have moral worth without manifesting a concern for doing what's right, but it also allows that an act can have moral worth in virtue of manifesting a concern (an instrumental concern) for doing what's right. This too is

advantageous, because such a concern is vital in dealing with moral uncertainty. Sometimes, we do not know what we should ultimately be concerned about. For instance, when it comes to promoting utility, we may not know whether we should have an ultimate moral concern only for promoting the overall sum of utility, or only for promoting the utility of each existing individual, or both. But if I were to know that I could trust someone to know that answer or at least trust that they know better than I do how to hedge one's bets as to which of these matter, then I should follow her advice out of non-instrumental concern for doing what's right. So, although there is something problematic about being motivated out a *non-instrumental* concern to do what's right given that this seems fetishistic, there's nothing problematic about being motivated by an *instrumental* concern to do what's right. And, so, we derive yet another important lesson.

Lesson 3: An act can have moral worth in virtue of its manifesting a non-instrumental concern for doing what's right.

And this means that we should reject rivals to the Concerns View that insist that an act can't have moral worth in virtue of its manifesting a non-instrumental concern for doing what's right—views such as those defended by Arpaly (2002).

Lastly, the Concerns View rightly implies that an act can have moral worth even if its agent believes that it's wrong. And the case that philosophers typically appeal to in support of this contention is Huck Finn. I'll begin, then, with Arpaly's helpful description of the case.

In Mark Twain's *The Adventures of Huckleberry Finn*, Huckleberry befriends Jim and helps him escape from slavery. While Huckleberry and Jim are together on their raft, Huckleberry is plagued by what he calls 'conscience'. He believes, as everyone in his society 'knows', that helping a slave escape amounts to stealing, and stealing is wrong. He also believes that one should be helpful and loyal to one's friends, but loyalty to friends is outweighed by some things, such as property rights, and does Miss Watson, Jim's owner, not have property rights? Hoping to find some excuse not to turn Jim in, Huckleberry deliberates. He is not very good at abstract deliberation, and it never occurs to him to doubt what his society considers common sense. Thus, he fails to find a loophole. 'What has poor Miss Watson done to me', he berates himself, 'that I can see her nigger go away and say nothing at all?' Having thus deliberated, Huckleberry resolves to turn Jim in, because it is 'the

right thing'. But along comes a perfect opportunity, and he finds himself psychologically unable to do it. (2002, p. 228)

Although this is perhaps as good a description based solely on Twain's novel as we can get, it leaves a lot unclear. First, it remains unclear how much concern Huck has for various things that matter morally. For instance, it's unclear whether Huck's level of concern for being a loyal friend is out of proportion to how much moral concern he should have for the welfare of others. Consequently, people like Johnson King worry that even if Jim were a serial killer on the run rather than a fugitive slave, Huck would still find himself psychologically unable to turn him into authorities (2020, p. 95). Second, it's not clear whether Huck cares about what's right or just about what's 'right' in the inverted comma sense (Markovits 2010, p. 208, n. 17). That is, it's unclear whether he thinks that turning Jim is what's genuinely right or is just what people in his society call 'right'. Third, this description doesn't tell us what sort of cases would count as relevantly similar to this one in assessing whether Huck's act of helping Jim was merely accidentally right. Are the relevantly similar cases those in which Huck again recognizes something of moral worth that the other members of his society fail to recognize as having moral worth, or are they those cases in which Huck recognizes that someone has a genuine property right and ensuring that that property is returned to its rightful owner requires disloyalty to a friend. On the Concerns View, whether Huck's act counts as merely accidentally right depends on what the context of our assessment is and what we take to be the relevantly similar cases in that context.

So, based on the above description alone, I don't think that we can adequately assess whether Huck's act of helping Jim elude the authorities is one that has moral worth. I suggest, therefore, that we fill in the details ourselves. So, let's assume that what motivates Huck to help Jim to elude the authorities is that Huck has come to see Jim as a fellow human being, and, consequently, he can't stand the thought of Jim's losing his rightful freedom. And let's assume that what justifies Huck in helping Jim is the fact that Jim is a fellow human being who would lose his rightful freedom if caught by the authorities. What's more, let's assume that Huck cares, and cares in the right proportions, for all other relevant moral matters. Thus, he cares more about protecting other people's welfare from serial killers than about being loyal to a friend.

And although he cares about people's property rights being respected, he rightfully assesses that helping an autonomous being go free is more morally important than ensuring that Miss Watson has what many in her society see as her property returned to her. Assume also that Huck has an instrumental concern for doing what's right, but that his non-instrumental concern for Jim and his freedom is rightfully greater than his instrumental concern for doing what's right. Given these assumptions, the Concerns View implies that Huck's act has moral worth. And this brings us to our fourth and final lesson.

Lesson 4: A right act can have moral worth even if its agent believes that it's wrong.

Of course, given that Huck believes that what he was doing was wrong, he clearly wasn't motivated to help Jim out of a concern for doing what's right. And this shows that we should reject rivals of the Concerns View that insist that an act can have moral worth only if its agent doesn't believe it to be wrong—views such as those defended by Singh (Forthcoming) and Sliwa (2016).

So, it seems that Concerns View has clear advantages not only over both the Kantian View and the Humean View but also over all the contemporary descendants of those two views.

5. Conclusion

As I understand the notion, what it means for a right act to have moral worth is for it to manifest a set of pertinent concerns that would never lead its agent to act impermissibly in a situation in which she's relevantly informed. And, of course, whether a set of concerns has this feature will just depend on whether it's appropriate in a certain sense. This much is trivial. Thus, to say, as the Concerns View does, that a right act has moral worth if and only if it manifests an appropriate set of concerns is trivial. But the Concerns View, as developed here, goes well beyond this trivial statement. It tells us that the appropriate concerns must include all pertinent concerns, that which concerns are pertinent is determined by the context, and that the qualitatively and quantitatively appropriateness of such concerns is determined by what the

agent's ultimate moral concerns should be. And, so, the resulting account of moral worth—that is, the Concerns View—is far from trivial. Indeed, as we've seen, its implications are contrary to all the extant theories in the literature. Unlike the Humean View and its contemporary descendants (e.g., Markovits 2010), it allows that a right act can lack moral worth even if it was performed for the right reasons. And, unlike the Kantian View and its contemporary descendants (e.g., Johnson King 2020), it holds that an act can have moral worth even if it doesn't manifest a concern for doing what's right.¹⁹

References

- Arpaly, N. (2015). "Moral Worth and Normative Ethics." *Oxford Studies in Normative Ethics* 5: 86–105.
- Arpaly, N. (2002). "Moral Worth." *The Journal of Philosophy*. 99: 223–45.
- Arpaly, N. and T. Schroeder (2014). *In Praise of Desire*. Oxford: Oxford University Press.
- Bradley, B. (2009). *Well-Being and Death*. Oxford: Clarendon Press.
- Darwall, S. (2010). "But It Would Be Wrong." *Social Philosophy and Policy* 27: 135–157.
- Furtak, R. A., (2018). *Knowing Emotions*. Oxford: Oxford University Press.
- Herman, B. (1981). "On the Value of Acting from the Motive of Duty." *Philosophical Review* 90: 359–382.
- Hooker, B. (2000). *Ideal Code, Real World*. Oxford: Oxford University Press.
- Howard, N. R. (Forthcoming). "The Goals of Moral Worth." *Oxford Studies in Metaethics*.

¹⁹ For helpful comments and criticisms on earlier drafts, I thank.... And I owe special thanks to Daniel Star for several very helpful conversations on moral worth when we were both residing in Princeton, NJ, in 2016.

- Hume, D. (2007). *A Treatise of Human Nature: A Critical Edition*. D. F. Norton and M. J. Norton (eds.). Oxford: Clarendon Press. [Originally published in 1739–1740.]
- Isserow, J. (Forthcoming). “Moral Worth: Having it Both Ways.” *The Journal of Philosophy*.
- Johnson King, Z. A. (2020). “Accidentally Doing the Right Thing.” *Philosophy and Phenomenological Research* 100: 186–206.
- Kagan, S. and P. Vallentyne (1997). “Infinite Value and Finitely Additive Value Theory.” *The Journal of Philosophy* 94: 5–26.
- Kant, I. (1998). *Groundwork of the Metaphysics of Morals*. Trans. M. Gregor. Cambridge: Cambridge University Press. [Originally published in 1785.]
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Markovits, J. (2010). “Acting for the Right Reasons.” *Philosophical Review* 119: 201–42.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Portmore, D. W. (Manuscript). “The Right, the Good, and Our Ultimate Moral Concerns.”
- Schaffer, J. (2000). “Trumping Preemption.” *The Journal of Philosophy* 97: 165–81.
- Schroeder, M. (2008). “Value Theory.” *The Stanford Encyclopedia of Philosophy* (Spring 2008), Edward N. Zalta (ed.), retrieved from
<URL><http://plato.stanford.edu/archives/spr2008/entries/value-theory/></URL>.
- Singh, K. (Forthcoming). “Moral Worth, Credit, and Non-Accidentality.” *Oxford Studies in Normative Ethics*.
- Sliwa, P. (2016). “Moral Worth and Moral Knowledge.” *Philosophy and Phenomenological Research* 93: 393–418.
- Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.
- Smith, M. (1987). “The Humean Theory of Motivation.” *Mind* 96: 36–61

Star, S. (Forthcoming). "Moral Worth, Normative Ethics, and Moral Ignorance." *Philosophical Analysis* (哲学分析).

Stocker, M. (1976). "The Schizophrenia of Modern Ethical Theories." *The Journal of Philosophy* 73: 453-466.

Suikkanen, J. (2008). "Consequentialism, Constraints, and Good-Relative-to." *Journal of Ethics and Social Philosophy* 3:1.

Williams, B. (1981). "Persons, Character and Morality." In *Moral Luck*, B. Williams, 1-18. Cambridge: Cambridge University Press.