

Moral Worth and Our Ultimate Moral Concerns

[Douglas W. Portmore](#)

Arizona State University
dwpportmore@gmail.com

Abstract: Some right acts have what philosophers call *moral worth*. A right act has moral worth if and only if its agent deserves credit for having acted rightly given that it was motivated in a way that made her acting rightly non-accidental. Given this, I argue that a right act has moral worth if and only if it issues from an appropriate set of concerns, where the appropriateness of these concerns is a function of what the agent's 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 A morally worthy act manifests the agent's virtuous motives such that she deserves credit for having acted rightly.² And, for an agent to deserve credit for having acted rightly, not only must her act be right, but her motives must be such that her acting rightly was no accident. More precisely, then, a right act has moral worth if and only if the agent deserves credit for acting rightly given that her act issued from a set of motives/concerns that made her acting rightly non-accidental.³

¹ By 'right', I mean 'objectively morally right'. And I use the term broadly to include both obligatory acts and supererogatory acts.

² An act manifests a given motive/concern if and only if we must appeal to that motive/concern in explaining its performance. Also, as I see it, to act rightly is merely to do what is right; it is not necessarily to do what is right *for the reasons that make it right* (cf. Hanser 2005).

³ I take 'moral worth' to be a term of art, and this, it seems to me, is how the term is most often used in the literature. Consider the following representative quotes: (1) "what matters [for moral worth] is that the action is in accord with duty and *it is no accident that it is*" (Baron 1995, p.131, emphasis in the original); (2) "when we say that an action has

As Immanuel Kant pointed out, not all right acts have moral worth. Take his example of the shopkeeper (G 4:397).⁴ The shopkeeper deals honestly with his customers, always giving them the correct change, but only because it's good for business. So, although he acts rightly in dealing honestly with his customers, his acts lack 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.⁵ Indeed, 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 issues from 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

moral worth, we mean to indicate (at the very least) that the agent acted dutifully from an interest...that therefore makes its being a right action the nonaccidental effect of the agent's concern" (Herman 1981, p. 366); (3) "for an action to have moral worth, it must not be a case of someone's merely *accidentally doing the right thing*" (Johnson King 2020, p. 191, emphasis in the original); (4) "because morally worthy action consists in the agent's deserving credit for doing the right thing, one cannot perform a morally worthy action by accidentally doing the right thing" (Singh **Forthcoming**); (5) "whether an act is only accidentally right affects whether it has moral worth, [because]...moral worth is a property of an agent's *action* only when they could not have easily acted wrongly" (Howard 2019, pp. 4 & 9, emphasis in the original); and (6) "morally worthy actions are motivated in a way that makes their rightness neither 'contingent' nor 'precarious'—they are...motivated in a way that makes them non-accidentally right" (Sliwa 2016, pp. 394 & 398). And, if this is the correct interpretation of this term of art, then a right act either does or doesn't have moral worth, depending on whether its agent does or doesn't deserve credit for having acted rightly given that her motives were or weren't such that her acting rightly was non-accidental. Nonetheless, I concede that, perhaps, not all authors have the exact same concept in mind when they write about 'moral worth'—for more on this, see Johnson King (2020, p. 189). Indeed, some authors hold that moral worth comes in degrees—see, e.g., Markovits (2010, p. 237). But even if we hold, as I do, that a right act has moral worth if and only if it issues from a set of concerns that makes it non-accidentally right and that whether an act issues from such a set is not something that comes in degrees, we could easily supplement my account to accommodate the intuition that there is something in the neighborhood of moral worth that comes in degrees. We could, for instance, add to my account the view that the degree to which an agent deserves praise for the motives that led her to act rightly depends on just how exceptional those motives were.

⁴ The 'G' stands for Kant's *Groundwork*, and the citation is given by volume and page number.

⁵ Thus, I'm assuming that one can act rightly (e.g., in accordance with duty) accidentally, such as where Kant's shopkeeper provides the correct change only because this just happens to coincide with promoting his self-interest. Of course, someone could object that the shopkeeper's duty is to provide the correct change out of respect for his customer's humanity and that this is a duty that he cannot fulfill accidentally given that it's a duty to perform an action from a good motive. But note that his having this duty entails his having a duty to provide the correct change, and this is a duty that he can fulfill accidentally—that is, from a motive that doesn't guarantee that he acts rightly.

miser might do something generous on occasion.⁶ And if, on that occasion, what moves her is an appropriate set of concerns, her act will have moral worth despite its 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 occasion in which everything goes her way: she finds the 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 finally getting the promotion for which she's long overdue. 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 stems from 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 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 stems from a virtuous set of concerns.

Also, as I see it, an act's having moral worth isn't the same as its agent being praiseworthy for performing it. Admittedly, an agent wouldn't deserve praise for performing an act if that act didn't have moral worth, but an act can have moral worth without its agent deserving praise for performing it. For what makes an act have moral worth is just that it was motivated in a way that made it non-accidentally right (see, e.g., Sliwa 2016, p. 398). But if its

⁶ I'm assuming that for an act to be generous it needn't issue from a stable disposition. So, I reject what Thomas Hurka (2006) calls *the dispositional view*, which identifies virtuous acts as those that issue from virtuous (and, consequently, stable) dispositions.

agent isn't responsible for having those motives (if, say, they were simply generated by a manipulator's implanted neurochip), then she won't deserve any praise for the acts that issued from them. So, in this respect, I differ from those who, like Nomy Arpaly, hold that "the moral worth of an action is the extent to which the agent deserves moral praise or blame for performing the action" (2002, p. 224). Unlike these people, I will not "speak interchangeably of a morally praiseworthy action and an action which has positive moral worth" (2002, p. 224). For whereas an act has positive moral worth if and only if it was motivated in a way that made it non-accidentally right, an act is morally praiseworthy if and only if its agent is, in virtue of having performed it, the deserving target of reactive attitudes such as pride, praise, gratitude, and admiration. And these two can come apart.

When it comes to an act's moral worth, it's not just the specific concerns that moved the agent to perform it that matter; her other concerns also matter. 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 reason. 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.⁷ 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

⁷ 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 a range counterfactual situations.

which she has the option of shooting the boy.⁸

The lesson, I take it, is that we must look not only at the specific motive/concern that moved her to act as she did, but also at her other pertinent concerns. And I'll be assuming that a lack of concern for something counts 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, her acting rightly was merely accidental, as it was fortuitous 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.

What counts as her pertinent concerns? 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's part of a set of concerns that would have led her to act wrongly in the relevantly similar situation in which she had the option of shooting the boy. In that situation, she would be led by her given set of pertinent concerns to act wrongly. Thus, her acting rightly in the given situation was merely fortuitous. Likewise, if Yunn's set of pertinent concerns had included a concern for the boy but not for dogs with spots, her acting rightly would have counted as merely accidental. For such a set would have led her to refrain from acting rightly (that is, to refrain from blocking the boy's blow) in the relevantly similar situation

⁸ Julia Markovits (2010, p. 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'll concede this point to Markovits provided 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 on the view that I'll be defending.

in which the boy was about to kick a Dalmatian rather than a poodle. So, again, her acting rightly would have been merely fortuitous in that it was just good luck that the dog in question happened to be a poodle rather than a Dalmatian.

Of course, not every concern that determines whether an agent would act rightly in some other situation is pertinent. 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 and how to protect a dog from physical violence. That said, I don't believe that there's any simple way of spelling out what the relevantly similar situations are. Consider, for instance, that we may wonder 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*. That is, we may wonder whether Yunn's having a concern for all mammals but not for any reptiles would make her acting rightly in *The Dog-Lover* count as merely accidental.

I suspect that the answer to such a question depends on the context in which it's being asked and on what's taken to be relevantly similar in that context. Thus, we could imagine one context in which we're concerned with whether Yunn is speciesist and, thus, with whether she has a concern for the welfare of all sentient beings and not just for the welfare of her fellow mammals. In that context, a situation in which the boy is intending to beat a snake would count as relevantly similar. And, so, her acting rightly in *The Dog-Lover* would count as merely accidental in that context. But we could also imagine a context in which we're merely concerned with whether Yunn is sufficiently altruistic with respect to the other members of her community and, thus, with whether she's willing to sacrifice her own welfare for theirs when appropriate. (And let's assume that her fellow humans and their canine companions—but no reptiles—count as members of her community.) In this context, the situation involving the snake wouldn't count as relevantly similar. And, so, her acting rightly in *The Dog-Lover* would count as non-

accidental in this context—we’re assuming that, in this case, she has a concern for all mammals and, thus, has the appropriate concern for the boy’s welfare.

As I see it, this context-sensitivity is a feature rather than a bug. For one, moral worth has to do with whether or not an agent was, given her motives, just lucky to have acted rightly, and, in general, whether someone counts as lucky is context sensitive. Take, for instance, the person who becomes infected with the Ebola virus and survives. On the one hand, we might consider her to be quite 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 exceedingly rare. For another, this sort of context-sensitivity helps us to explain why there tends to be so much intractable disagreement about some of the cases discussed in the literature on moral worth—such as the much-discussed case of Huck Finn. 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 relevantly similar 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.⁹ So, in the rest of the paper, I will just focus on cases in which I suspect that there will be wide agreement as to what the relevant similarity relation is.

As we’ve just seen, what determines whether an act has moral worth is whether it issues from a set of motives/concerns that makes the agent’s having acted rightly non-accidental.¹⁰ And, as we learned from the shopkeeper case, an act counts as merely accidentally right if it issues from a selfish motive given that such a motive would lead one to act wrongly in the relevantly similar situation in which it’s in one’s self-interest to act wrongly.¹¹ And, as we

⁹ 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).

¹⁰ There is wide consensus on this point—see those quoted in [note 3](#) as well as Arpaly (2002, p. 225), Isserow (2020, p. 532), and Markovits (2010, pp. 206 & 211). Zoë Johnson King (2020, p. 192) agrees that there’s consensus on this point. Nonetheless, for some rare skepticism on this point, see Rozeboom (2017, p. 5, n. 14).

¹¹ 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 coincide with promoting her self-interest. And this, I think, is why there seems to

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 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 of protecting the dog 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 context. Lastly, 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 accept the following criterion for moral worth.

The Non-Accidentality Criterion: A right act has moral worth if and only if it issues from a set of pertinent concerns that would never lead its agent to act impermissibly in a situation in which she's relevantly informed.¹²

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 maximizing profits.

¹² 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 stagehand with such a concern would be led to determine whether it's loaded with blanks or live ammunition. Of course, if the stagehand 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 objective (i.e., fact-relative) sense of 'ought', I'll be concerned with agents who are *relevantly* informed as opposed to *fully* informed.

I'm assuming that the relevant sort of non-accidentality should be understood in terms of a *modal connection* between the agent's concerns and her acting rightly, but see Singh (Forthcoming) for a very different approach, one where non-accidentality is instead understood in terms of an *explanatory connection* between the agent's concerns and her acting rightly. Also, see Johnson King who thinks that non-accidentality should be understood in terms of deliberate agency, such that one does something non-accidentally if and only if she does that something deliberately (2020, p. 202). Now, as I'll argue below, there are cases where the views of both Singh and Johnson King imply that an act that is intuitively non-accidentally right is merely accidentally right as well as cases where their views imply that an act that it intuitively merely accidentally right is non-accidentally right. This, I believe, vindicates my choice to interpret non-accidentality in terms of a modal connection as opposed to either an explanatory connection (as Singh does) or a connection in terms of deliberate agency (as Johnson King does). Also, the sorts of examples that

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 explain at the end of the next section (section 2) what this fundamental flaw is. In section 3, I continue to make my case against these two simple views by presenting some additional counterexamples to them. And this leads me to introduce a new concept—the concept of an ultimate moral concern—in section 4. In section 5, I employ this concept in developing a new account of moral worth and show how this account compares favorably to its rivals. Lastly, in section 6, I conclude by explaining the extent to which this new account of moral worth is and isn't trivially true.

2. The Simple Kantian View and the Simple 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 spawned. 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).

The Simple 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

Johnson King gives (see 2020, p. 194) in support of her claim that, for any act-type *A*, an agent has non-accidentally *A*-ed if and only if she deliberately *A*-ed don't give us any reason to accept it over my contrary claim that, for any act-type *A*, an agent has non-accidentally *A*-ed if and only if the pertinent concerns that resulted in her *A*-ing guarantee that she will *A* in any situation in which she is relevantly informed.

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.

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 views, not because I think that it's correct, but because its simplicity allows me to more perspicuously illustrate how they differ from each other.

There are some clear counterexamples to the Simple Kantian View. Here's one.

The Empathic: A man named Christoforos denies that chimpanzees are morally considerable beings, for he mistakenly believes 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.¹³ So, when the opportunity arises, he's moved to help the chimpanzee escape back into the wild out of a non-instrumental concern to alleviate what he correctly perceives to be the chimpanzee's distress.¹⁴ And this is his

¹³ 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. 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, the propositions that he's willing to assent to), he inaccurately represents the world as being one in which the chimpanzee is not suffering. Fortunately, his actions are being guided by what his emotions are telling him rather than by what his avowed beliefs are telling him.

¹⁴ He wants to alleviate the chimpanzee's distress for its own sake and not merely as a means to alleviating the unpleasantness that the chimpanzee's distress is causing him. Thus, assuming the Humean theory of motivation, 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. 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, 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. He does this knowing that it involves substantial risk to himself.¹⁵

On the Simple 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 issued from an appropriate set of concerns—one that made his acting rightly non-accidental. After all, he had a concern not only to alleviate the chimpanzee's distress, but also to safeguard the welfare of both the poachers and the authorities. He even had a concern for doing what's right. What's more, the magnitude of each of these concerns was, we'll assume, 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 Simple Kantian View implies, his act of helping the chimpanzee was non-accidentally right and, consequently, has moral worth. Thus, the Simple Kantian View conflicts with the accidentality criterion and should, therefore, be rejected. It conflicts with the accidentality criterion, because although having a non-instrumental concern for doing what's right (whatever that may be) will

motivating reason consists in both a non-instrumental desire to alleviate the chimpanzee's distress and the belief that were he to help the chimpanzee escape back into the wild he would thereby do so.

¹⁵ More commonly, philosophers (such as Arpaly 2002) cite the case of Huck Finn as a putative counterexample to the Simple 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 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.

unerringly lead one to do what's right in any situation in which one is relevantly informed, a set of concerns that doesn't include a non-instrumental concern for doing what's right (whatever that may be) can also unerringly lead one to do what's right in any situation in which one is relevantly informed, as *The Empathetic* shows.

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).¹⁶ Here's a simple version of his view.

The Simple Humean View: A right act has moral worth if and only if it manifests a non-instrumental concern for performing acts with 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 with RMF and the belief that were she to perform this act she would perform an act with RMF.

This view is also subject to counterexample, though what sort of counterexample it's subject to depends on which of the following two versions of it we have in mind.

On the fundamentalist version, 'RMF' refers to whatever the *fundamental* right-making feature of acts is. Thus, if maximizing act-utilitarianism is correct, 'right-making feature RMF' refers to 'the feature of maximizing utility'.¹⁷ Given this view and the Humean theory of

¹⁶ The 'T' refers to Hume's *A Treatise of Human Nature*, and the citation is given by book, part, section, and paragraph number.

¹⁷ Nomy Arpaly and Timothy Schroeder seem to endorse the fundamentalist version of the Simple 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, all caps in the original). Thus, as Arpaly explains in another work, "a morally worthy action stems from a commitment to the right and the good *correctly*

motivation, a right 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 utility and the belief that, were she to perform this act, she would maximize utility.

On the non-fundamentalist version, by contrast, 'RMF' can refer to any right-making feature, fundamental or non-fundamental. To illustrate, assume that maximizing act-utilitarianism is correct and suppose that I would maximize 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 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 Simple 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 utility and the belief that, were I to perform this act, I would maximize 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 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 Simple Humean View, 'right-making feature RMF' can refer to 'the feature of maximizing utility', 'the feature of saving many lives', or any other right-making feature.

Unfortunately, both versions of the Simple 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 Simple Humean View implies 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

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).

maximizing act-utilitarianism is correct, what fundamentally makes it right to block the boy's blow is that doing so would maximize utility.¹⁸ 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 utility under the circumstances. So, on the non-fundamentalist version, Yunn's act has moral worth. But, as we saw above, Yunn's act was merely accidentally right given that she had zero concern for the boy's welfare. And, thus, she would have shot the boy had that been an option. Thus, the non-fundamentalist version of the Simple Humean View conflicts with the non-accidentality criterion and should, therefore, be rejected.

The reason the non-fundamentalist version of the Simple Humean View conflicts with 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 confer moral worth on your act are ones that could lead you to act wrongly in other situations in which you are relevantly informed, making your acting rightly in this instance merely accidental. 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 Simple Humean View, we would have to adopt the fundamentalist version of the Simple 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 Simple Humean View is also unacceptable. The problem is that it makes moral worth too hard to come by. As Daniel Star (Forthcoming) has

¹⁸ 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 her if blocked by her.

pointed out, people rarely conceptualize their actions as meeting some fundamental moral criterion.¹⁹ Even self-professed maximizing act-utilitarians rarely conceptualize what they're doing as maximizing 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 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 maximize utility and the belief that by acting this way she would maximize utility (or, substitute here whatever the correct 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.

In any case, *The Empathic* seems to be a clear counterexample to the fundamentalist version of the Simple Humean View. Christoforos wasn't motivated out a non-instrumental concern for anything such as maximizing utility, abiding by the ideal code of rules, or acting in accordance with the categorical imperative. Rather, he was motivated simply out of a non-instrumental desire to alleviate what he correctly perceived to be the chimpanzee's 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, say, doing what would maximize 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 utility (and, thus, to act rightly) in any situation in which he is relevantly informed. Therefore, we should reject the fundamentalist version of the Simple Humean View, for it too conflicts with the non-accidentality criterion.

So, we should reject both the Simple Kantian View and the Simple Humean View. We should reject the Simple Kantian View because it gets the wrong verdict in *The Empathic*. And we should reject the Simple Humean View because it gets the wrong verdict in either *The Empathic* or *The Dog-Lover*, depending on whether we're considering its fundamentalist or non-

¹⁹ See Star (Forthcoming). See also Howard (Forthcoming).

fundamentalist version.²⁰ 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 issue from an appropriate set of concerns have moral worth.

More specifically, the Simple Kantian View 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 should not care about doing what's right for its own sake. Rather, we should care about doing what's right only as a means to ensuring that we act as we would naturally be inclined to act if we were to have all and only the appropriate concerns. That is, we should care about acting morally, not for its own sake, but only for the sake of the ends at which morality properly aims. 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 Simple 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 don't know what's right because either we 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

²⁰ Jessica Isserow 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" (2020, p. 550). I believe that her view will suffer the same fate as the Simple 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*.

is—as a result of, say, the reliable testimony of someone whom we can trust to know what’s right in this situation. And, in those instances, we should care about doing what’s right as a means to doing right by that which 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 fail in the moment to care directly 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.²¹ Yet, according to the Simple 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 Simple Kantian View and the Simple Humean View fail to accommodate the plausible idea that all and only those right acts that issue from 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 Simple Kantian View fails in requiring us to have an inappropriate concern (specifically, a non-instrumental concern for doing what’s right), and the Simple Humean View fails in prohibiting us from having an

²¹ Hallvard Lillehammer has provided a few nice examples, including this one: “Consider someone who goes to a party during a phase when she is tired of her husband. At the party she meets a very charming person and is tempted to have an affair. She judges that it would be wrong to have an affair on account of her husband’s feelings. But she is temporarily indifferent to her husband’s feelings. However, she has a standing *de dicto* desire to do what is right which, together with her moral judgement, causes her to do the right thing, in spite of the absence of a *de re* desire to do the right thing and the presence of a *de re* desire to do the wrong thing. If there is anything in this case which prevents this person from being good it is not her standing desire to do what is right, where this is read *de dicto*. For this desire is playing the role of an internalised norm that prevents her from being tempted to do wrong” (1997, p. 192). Of course, it would be morally preferable for her to care about her husband’s feelings even on this occasion, but, absent that, there’s nothing morally objectionable about her having a *de dicto* desire to do what’s right as a means to ensuring that she acts in the way that she would naturally be inclined to act if she had had all the appropriate *de re* desires. After all, a *de dicto* desire to do what’s right will unerringly lead her to act rightly in any situation in which she is relevantly informed. Thus, the problem with the Simple Kantian View is that it fails to recognize that, although having a *de dicto* desire to do what’s right is sufficient to ensure that one acts rightly whenever one is relevantly informed, it isn’t necessary.

appropriate concern (specifically, an instrumental concern for doing what's right). Of course, in making these arguments, I've relied heavily on the notion of what an agent should ultimately be concerned about. Unfortunately, this notion has been undertheorized. So, I will need to take a brief digression from our discussion of moral worth to explicate it. But, before I do, I want to present two different sorts of counterexamples, one to the Simple Kantian View and other to the Simple Humean View.

3. *Two Other Counterexamples*

I find that some are not convinced by *The Dog-Lover* and *The Empathic* to reject both the Simple Kantian View and the Simple Humean View. So, let me offer two other rather different sorts of counterexamples.²² They are as follows:

Golfing for Rightness' Sake: Unless a patient's tumor is removed this afternoon, he'll die this evening—though not painfully. Cutting and stitching this afternoon by the only two available doctors, Slice and Patch, is the only thing that can save him. For Slice is the only one who can cut out the tumor, and Patch is the only one who can stitch him up afterwards. If there is either cutting without stitching or stitching without cutting, the patient's death will be physically agonizing. It would even be cruel for one of the two doctors to show up to the hospital knowing that the other won't, as this would only needlessly get the patient's hopes up, making his death psychologically agonizing. Unfortunately, both Slice and Patch want their patient to die. Consequently, each has decided not to take part in the operation regardless of what she thinks the other might do. And each is immovable in this regard. What's more, each knows this about the other. Thus, each knows that, given the other's unwillingness to do her part in saving the patient, taking her own husband golfing is the right thing to do. For let's assume that,

²² These are adapted from David Estlund's case entitled "Slice and Patch Go Golfing" (2017: 53), which in turn was adapted from Donald Regan's case of Whiff and Poof—see his 1980. But neither used this sort of example to test our intuitions about moral worth.

given the unwillingness of the other to participate in the operation, the best thing that she can do is to take her own husband golfing. Now, neither doctor has any concern for anyone but herself, but each has a non-instrumental concern for doing what's right (whatever that may be). Unfortunately, though, this desire is not as strong as each's desire that the patient dies, which is what explains why neither is willing to do her part in saving the patient. So, in the end, each doctor takes her husband golfing out of a non-instrumental concern for doing what's right. Consequently, they together produce what I've labelled in [Table 1](#) as outcome O₄—the outcome in which each goes golfing with her own husband while the patient dies painlessly.

Golfing for RMF's Sake: This case is exactly the same as the above except that, in this case, each doctor has a non-instrumental concern for performing acts with right-making feature RMF rather than a non-instrumental concern for doing what's right. And, given the other's unwillingness to do her part in saving the patient, the right thing for each doctor to do is to take her own husband golfing. For let's assume that each doctor has promised to take her husband golfing this afternoon. Thus, no matter what the correct moral theory is, taking one's own husband golfing will have right-making feature RMF.

[INSERT TABLE 1 HERE – IT CURRENTLY APPEARS [AFTER THE REFERENCES](#)]

I'm assuming that, on the correct moral theory, one is required to keep one's promises whenever there's nothing better to do.²³ And there is nothing better for either doctor to do. For, given Patch's unwillingness to cooperate, Slice's showing up to the hospital has no chance of saving the patient and every chance of making the patient's death psychologically agonizing. Likewise, given Slice's unwillingness to cooperate, Patch's showing up to the hospital has no chance of saving the patient and every chance of making the patient's death psychologically agonizing. So, given the circumstances, each doctor should stay away from the hospital and

²³ I'm also assuming that, on the correct moral theory, it would be permissible to break one's promise to take one's spouse golfing in order to save a life.

instead take her husband golfing, which is the best thing that either of them can do under the circumstances. Of course, many non-maximizing theories would deny that an agent must always perform her best option. But every plausible moral theory (maximizing or non-maximizing) will require an agent to perform her best option when doing so would fulfill a promise, maximize utility, and neither harm nor disrespect anyone. So, it seems safe to assume that, on any plausible moral theory, the right thing for each doctor to do is to take her own husband golfing. What's more, taking her husband golfing will have RMF. For whether 'RMF' stands for 'maximizing utility', 'keeping one's promise when there's nothing better to do', 'respecting humanity always as end and never as a mere means', or what have you, taking her husband golfing will have this feature.

But whether we're talking about *Golfing for Rightness' Sake* or *Golfing for RMF's Sake*, these doctors don't deserve credit for acting rightly. After all, their motives were such that each would take her husband golfing regardless of what the other was willing to do and, thus, regardless of whether it would be wrong to do so. Thus, each just lucked out in that the other doctor was unwilling to do her part in saving the patient. For even had the other doctor been willing to do her part, she would have still taken her husband golfing, which, in those circumstances, would have been wrong. It seems, then, that neither doctor's act of taking of her husband golfing counts as something non-accidentally right. Thus, it seems that neither has moral worth. Yet, both the Simple Kantian View and the Simple Humean View imply otherwise.

In *Golfing for Rightness' Sake*, each doctor was motivated to take her husband golfing out of a non-instrumental desire to do the right thing and the belief that this was the right thing to do. So, on the Simple Kantian View, each doctor's act of taking her husband golfing had moral worth. And, in *Golfing for RMF's Sake*, each doctor was motivated to take her husband golfing out of a non-instrumental desire to perform an act with RMF (e.g., with the feature of being a promise-keeping) and the belief that doing so has RMF. So, on the Simple Humean View too, each doctor's act of taking her husband golfing had moral worth. But although each doctor did the right thing in taking her husband golfing, neither of them deserve credit for acting rightly. Again, it was merely fortuitous that they acted rightly. Indeed, the same motives that led them

to act rightly in this situation would have led them to act wrongly in the relevantly similar situation in which the other doctor was willing to do her part in saving the patient.

At this point, I think that we can be confident that both the Simple Kantian View and the Simple Humean View are unacceptable. Both are subject to counterexample. And both are fundamentally flawed in that they both fail to accommodate the plausible idea that all and only those right acts that issue from an appropriate set of concerns have moral worth. And this is what explains why these views get the wrong verdicts in the above cases. For had Slice and Patch had the appropriate concerns, they would each have the appropriate concern for their patient. And had they each had the appropriate concern for their patient, they would have worked together to save him. For I'm assuming that what's appropriate is that they care more about saving their patient than about keeping their relatively trivial promises to take their husbands golfing. But what makes this appropriate? And what in general determines whether a given concern is appropriate? I now take up this issue.

4. Our Ultimate Moral Concerns

Unfortunately, most moral theories don't tell us what we morally ought to have a concern for, let alone how much concern we should have for those things. If they did, they would need to tell us what our ultimate moral concerns should be—our ultimate moral concerns being those that don't derive from any other more fundamental moral concern and are, therefore, the concerns from which all our derivative moral concerns derive. 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.

I believe that there are at least three reasons to think that a complete moral theory owes us an account of what our ultimate moral concerns should be. First, 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. Second, we need to know what our ultimate moral concerns should be so that we can determine whether our concerns are

appropriate and, thus, whether the acts that issue from them have moral worth. And, third, 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 determines whether it 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 it prescribes for them. And this is important, because we should reject any theory that's incoherent in this sense.

To see why we should reject any such theory, 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 shown, 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 even though 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 it gives them. 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.²⁴ 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 coherent or not. Now, admittedly, many 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 decision procedure. There are exceptions, of course. Some (particularly

²⁴ 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.

virtue ethicists) have offered a theory of virtue, and, from such a theory, we can deduce what our ultimate concerns should be.²⁵ Nevertheless, the focus has tended to be on providing a criterion of rightness, especially among deontologists and consequentialists. Indeed, the only deontologist or consequentialist that I've known to have explicitly specified what an agents' ultimate moral concerns should be is Brad Hooker. But this is clearly a mistake given that it's not just rule-consequentialism that's potentially subject to an incoherence objection. All moral 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 them with 2 hedons 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 you. 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.²⁶

Of course, some may doubt whether I've actually identified a new and important notion. For some may suspect that this notion of an ultimate moral concern is just equivalent to some

²⁵ Motive utilitarianism is also an exception—see Adams 1976.

²⁶ 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 is impartially defensible just when there is no alternative that would produce a greater sum of utility than it would.

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 (1984) calls "our theory-given aims." But, on Parfit's view, our theory-given aims are provided by that 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). 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, which is what our theory-given aim is on rule-consequentialism. Rather, agents 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 right/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 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.²⁷ 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, for the ultimate moral concerns that an agent should have need not be for what's good, better,

²⁷ In the case of an incoherent theory (e.g., a version of rule-consequentialism that holds that everyone's ultimate moral concern should be to maximize utility), a relevantly informed agent could be led by the ultimate moral concerns that it prescribes for her to act in a way that that theory deems wrong. Take, for instance, the version of rule-consequentialism just mentioned: an agent would be led by an ultimate concern for maximizing utility to violate the ideal code when violating that code would maximize utility, which is something that this theory deems wrong. Nevertheless, the correct moral theory can't be incoherent. And this is why having the correct ultimate moral concerns will unerringly lead one to act rightly in any situation in which one is relevantly informed.

or best. To illustrate, suppose that I must choose either to save my own child or to enable some stranger to save her child. And suppose that the latter would be slightly better than the former. For, perhaps, goodness is simply a function of utility and there would be slightly more utility as a result of my choosing to enable the stranger to save her child. Even so, we might think that I should ultimately have a greater moral concern for the welfare of those near and dear to me such that I ought, morally speaking, to prefer saving my own child to enabling a stranger to save hers. So, the ultimate moral concerns that one should have needn't be dictated by what's good, better, or best. And, thus, the two notions are distinct.

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 rightness and goodness. Yet, as I've shown, this notion is crucial to moral theorizing, for we should reject theories that are incoherent, and whether a theory is incoherent just depends on its account of what our ultimate moral concerns should be. We've also seen that such an account is crucial to our determining whether the concerns that issue in an act are appropriate and, thus, the kind that confer moral worth on an act. For we need to know what the agent's ultimate moral concerns should be in order to determine whether the concerns that issue in her acts are appropriate. Thus, in the next section, I return to the issue of moral worth and show how this notion of an ultimate moral concern can help us to develop a plausible alternative to both the Simple Kantian View and the Simple Humean View.

5. 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 issues from an appropriate set of concerns—that is, a set that includes all and only pertinent concerns, each of which must be both qualitatively and quantitatively appropriate, which in turn is determined by 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 determined by what the agent’s ultimate moral concerns should be. To illustrate, suppose, as I argued above, that acting rightly is not something for which an agent should have an ultimate moral concern. In that case, she should be concerned with acting rightly only insofar as acting rightly is a means to furthering the ends for which she should have an ultimate moral concern. 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 the ends for which she should have an ultimate moral concern.

To take another example, imagine that an agent should have an ultimate moral concern for promoting each existing individual’s utility 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, given the appropriate ultimate moral concerns, 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 concerns from which her act stems, 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: specifically, a concern for the welfare of the boy. For it's plausible to suppose that, for each sentient being, Yunn should have an ultimate moral 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 in combination with other appropriate concerns lead her to act wrongly in various relevantly similar situations, such as the one in which she has the option of shooting the boy. And since Yunn's act of blocking the boy's blow doesn't stem from an appropriate set of concerns, the Concerns View rightly implies that it lacks moral worth.²⁸ And there's a lesson 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 reason (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 act 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 stem from an appropriate set of concerns. In notifying the authorities and hiding the poacher's 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, we're assuming. Likewise, his act of helping the chimpanzee escape stems from an appropriate set of concerns. Of course, it doesn't manifest a non-

²⁸ The same holds for Slice and Patch in both *Golfing for Rightness' Sake* and *Golfing for RMF's Sake*. On the Concerns View, each doctor's act of taking her husband golfing lacks moral worth given that each doctor lacked a pertinent concern: specifically, a concern for the patient's welfare.

instrumental concern for doing what's right given that he doesn't think that helping the chimpanzee escape is right; he thinks, rather, that it's morally neutral. Nevertheless, his act does manifest an empathetic concern for alleviating what he correctly perceives to be the chimpanzee's 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 distress. So, in this case too, the Concerns View gets the intuitive verdict. What's more, the Concerns View provides us with another important lesson.

Lesson 2: An act can have moral worth even if it doesn't manifest a non-instrumental 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 non-instrumental concern for doing what's right—views such as those defended by Herman (1981), Johnson King (2020), Singh (Forthcoming), and Sliwa (2016).

Unlike the Simple 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 someone 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, someone 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 (instrumental or non-instrumental) 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 both temptation and uncertainty. Sometimes, we're tempted to do what's wrong but are moved by our desire to do what's right to resist that temptation. Other times, we do not know (or are uncertain about) 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 for promoting the overall sum of utility, for promoting the utility of each existing individual, or for both. But if I were to know that I could trust someone to know the answer or at least trust them to know better than I do how to hedge one's bets with respect to the relevant possibilities, then I should follow her advice out of an 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 is fetishistic, there's nothing problematic about being motivated by an *instrumental* concern to do what's right where one faces temptation or moral uncertainty. And, so, we derive yet another important lesson.

Lesson 3: An act can have moral worth in virtue of its manifesting a concern (specifically, an 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 concern for doing what's right—views such as those defended by Arpaly (2002).

The Concerns View also has the advantage of implying 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 that of 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 as we can get that's based solely on Twain's novel, it leaves a lot unclear and unspecified. 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, 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 in is what's genuinely right or is just what people in his society call 'right'. Third, this description doesn't specify what sorts 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 responds emotionally to something of moral importance despite believing that it is not of moral importance? Or are they those in which Huck helps a lawbreaker (perhaps, even a serial killer) against his better moral judgment and out of loyalty to a friend? On the Concerns View, whether Huck's act counts as merely accidentally right (and, thus, has moral worth) 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 relevant details ourselves. So, let's assume that what motivates Huck to help Jim to elude the authorities is that Huck has come to recognize Jim as an autonomous being (that is, as a being with the capacity for rational self-rule), and is, consequently, concerned that Jim not lose his freedom. And let's assume that what morally justifies Huck in helping Jim is the fact that Jim is an autonomous being who would lose his

freedom if caught by the authorities. What's more, let's assume that Huck cares more about protecting other people's welfare from serial killers than about being loyal to a friend. And although he cares about people not breaking the law, he cares more about letting an innocent and autonomous being go free than with ensuring that no one breaks an unjust law. Assume also that Huck has an instrumental concern for doing what he thinks is right, but that his non-instrumental concern for Jim and his freedom is rightfully greater than his instrumental concern for doing what he thinks is right. Lastly, assume that the context is one in which we're concerned with whether Huck would aid a lawbreaker when, and only when, it's actually morally right for him to do so, not with whether he would do so even when he *believes* that that would be morally wrong to do so. Given these assumptions, the Concerns View implies that Huck's act has moral worth. And this brings us to our fourth lesson.

Lesson 4: A right act can have moral worth even if its agent believes that it's wrong (Markovits 2010, p.208).

Of course, given that Huck believes that it's wrong for him to help Jim allude the authorities, he clearly wasn't motivated to do so out of a concern for doing what's right. But given our assumption that, in this context, what we're interested in is whether Huck would ever aid a lawbreaker when it's actually morally wrong to do so, not with whether he would do so even when he *believes* that it would be morally wrong to do so, his lack of a strong concern for doing what he believes is right just isn't pertinent. And, so, this case 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 is wrong—views such as those defended by Johnson King (2020) and Sliwa (2016).

Lastly, cases such as *Golfing for Rightness' Sake* and *Golfing for RMF's Sake* show us that being motivated to do what's right, either because it's right or because it has right-making feature RMF, is insufficient to confer moral worth on an act. Consider that, in *Golfing for Rightness's Sake*, each doctor takes her husband golfing because it's the right thing to do, and, yet, their acts are only accidentally right given that they would have done the exact same thing even if it had been wrong to do so. So, deliberately acting rightly is insufficient to confer moral

worth on an act. And, in *Golfing for RMF's Sake*, each doctor takes her husband golfing because this has right-making feature RMF. But, in this case too, their acts are only accidentally right given that they would have done the exact same thing even if it hadn't had RMF. And this brings us to our fifth and final lesson.

Lesson 5: A right act can lack moral worth even if its agent performs it because it's right (and even if its agent performs it because it has right-making feature RMF).

So, we should reject rivals of the Concerns View that insist that an act will have moral worth if it is an instance of its agent deliberately acting rightly (or if it is an instance of its agent acting on the basis of sufficient moral reasons, conceived as such)—views such as those defended by Johnson King (2020) and Singh (Forthcoming).²⁹

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

6. Conclusion

As I understand the notion, what it means for a right act to have moral worth is for it to issue from 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 issues from an appropriate set of concerns is trivially true. 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 the pertinence of our concerns is determined by the context, and that the qualitative and quantitative appropriateness of such concerns is determined by what the

²⁹ Assume that, in *Golfing for Rightness' Sake*, each doctor deliberately acts rightly. And assume that, in *Golfing for RMF's Sake*, each doctor takes her husband golfing because this has right-making feature RMF, knowing that this is sufficient to make it right.

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, many of its implications are contrary to those proffered by others in the literature. Unlike the Simple 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 Simple 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.³⁰

References

- Adams, R. (1976). "Motive Utilitarianism." *The Journal of Philosophy* 73: 467–81.
- 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.
- Baron, M. (1995). *Kantian Ethics Almost without Apology*. Ithaca, NY: Cornell 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: 135–57.
- Estlund, D. (2017). "Prime justice." In *Political Utopias*, K. Vallier and M. Weber (eds.), 35–55. Oxford: Oxford University Press.
- Furtak, R. A., (2018). *Knowing Emotions*. Oxford: Oxford University Press.

³⁰ For helpful comments and discussions, I thank Cheshire Calhoun, Nathan Robert Howard, Paul Hurley, Jessica Isserow, Zoë Johnson King, Andrew Khoury, Aaron Rizzieri, Grant Rozeboom, and Keshav Singh. 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.

- Hanser, M. (2005). "Permissibility and Practical Inference." *Ethics* 115: 443–70.
- Herman, B. (1981). "On the Value of Acting from the Motive of Duty." *Philosophical Review* 90: 359–82.
- Hooker, B. (2000). *Ideal Code, Real World*. Oxford: Oxford University Press.
- Howard, N. R. (Forthcoming). "The Goals of Moral Worth." *Oxford Studies in Metaethics*.
- Howard, N. R. (2019). "One Desire Too Many." *Philosophy and Phenomenological Research*, Early view version 1–16. <https://doi.org/10.1111/phpr.12642>.
- 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.]
- Hurka, T. (2006). "Virtuous Act, Virtuous Dispositions." *Analysis* 66: 69–76.
- Isserow, J. (2020). "Moral Worth: Having it Both Ways." *The Journal of Philosophy* 117: 529–56.
- 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.
- Lillehammer, H. (1997). "Smith on Moral Fetishism." *Analysis* 57: 187–95.
- 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."
- Regan, D. (1980). *Utilitarianism and Co-operation*. New York: Oxford University Press.

- Rozeboom, G. J. (2017). "The Motives for Moral Credit." *Journal of Ethics and Social Philosophy* 11: 1–29.
- Schaffer, J. (2000). "Trumping Preemption." *The Journal of Philosophy* 97: 165–81.
- 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, D. (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: 453–66.
- Williams, B. (1981). "Persons, Character and Morality." In *Moral Luck*, B. Williams, 1–18. Cambridge: Cambridge University Press.

Table 1: Slice and Patch

	Slice cuts	Slice goes golfing with her husband
Patch stitches	They produce the best world that they could together produce: (O_1) the patient is saved, although both husbands are disappointed not to go golfing. Each maximizes utility.	They produce the world that's tied for the worst world that they could together produce: (O_2) the patient dies in agony, Patch's husband is disappointed not to go golfing, but Slice's husband is glad to go golfing. Neither maximizes utility.
Patch goes golfing with her husband	They produce the world that's tied for the worst world that they could together produce: (O_3) the patient dies in agony, Slice's husband is disappointed not to go golfing, but Patch's husband is glad to go golfing. Neither maximizes utility.	They produce the second-best world that they could together produce: (O_4) the patient dies painlessly and both husbands are glad to go golfing. Each maximizes utility.