

Moral Worth and Our Ultimate Moral Concerns

Douglas W. Portmore

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 they deserve credit for having acted rightly.² And, for an agent to deserve credit for having acted rightly, not only must their act be right, but their motives must be such that their acting rightly was no mere accident. More precisely, then, a right act has moral worth if and only if it was motivated in a way that makes it non-accidentally right.³

As Kant noted, not all right acts have moral worth. Take his example of the shopkeeper (G 4:397). 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

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

² 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 'Moral worth' is not an ordinary language term (see Johnson King 2020, 189). It is, instead, a term of art that's most often used to describe an act as being motivated in a way that makes it non-accidentally right. To illustrate, 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, 131); (2) "when we say that an action has 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, 366); and (3) "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, 394 & 398). See also Arpaly (2002, 225), Howard (2019, 4–9), Isserow (2020, 532), Johnson King (2020, 191), Markovits (2010, 206 & 211), and Singh ([Forthcoming](#)).

with his customers, his acts lack moral worth in that they were motivated solely out of a concern to maximize profits, a concern that would have led him to short-change his customers had they been less savvy. Thus, his acting rightly was merely fortuitous.

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 miser might do something generous on occasion.⁴ And if, on that occasion, what moves them is an appropriate set of concerns, then their act will have moral worth even if it's out of character. Likewise, someone with a good character might act from bad motives on occasion, and, when they do, their 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 them to act wrongly in other situations, an agent's character has to do with whether they're disposed to have the appropriate concerns in a wide range of situations, even if not necessarily in the given one.

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 acting rightly unless their acting rightly was non-accidental. But an agent can deserve praise for, say, *acting selflessly* even if they don't deserve praise for *acting rightly*. Thus, that an act was non-accidentally right isn't the only possible reason for praising it (see Johnson King 2020, 191). I deny, then, that "the moral worth of an action is the extent to which the agent deserves moral praise or blame for performing the action" (Arpaly 2002, 224). And, so, I will not "speak interchangeably of a morally praiseworthy action and an action which has positive moral worth" (Arpaly 2002, 224), for I see no reason to introduce a technical term as a stand-in for an ordinary one.

Also, unlike some others (e.g., Markovits 2010), I don't see moral worth as something that comes in degrees. For either an agent's motives were such that their acting rightly was non-accidental or they weren't. Nonetheless, I concede that some may use the term differently than I

⁴ I'm assuming that for an act to be generous it needn't issue from a stable disposition to be generous. So, I reject what Hurka (2006) calls *the dispositional view*.

do. But most in the literature use the term as I do. Of course, one may still worry that as we (the majority) use it, it refers to some rather uninteresting notion. But this is not the case. For there are at least two reasons to take an interest in whether an act was motivated in a way that makes it non-accidentally right. First, as pointed out above, whether we should praise someone for acting rightly depends on whether their actions were motivated in a way that makes them non-accidentally right. Second, we're often interested in whether someone who has acted rightly in one situation can be counted on to act rightly in other situations. And that's precisely what this notion tells us.

I also differ from several others (e.g., Markovits 2010 and Sliwa 2016) in thinking that the moral worth of an act doesn't just depend on the specific concerns that moved the agent to perform it; I believe that the agent's other concerns matter as well. To see why, consider the following.

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. 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 lead her to do the wrong thing in other relevantly similar situations, such as the one in which she has the option of shooting the boy.⁵

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 their fanatical love for dogs, they would not have saved these strangers had the choice been between saving them and saving their dog. Markovits claims that, assuming that the dog-lover's preference for saving their dog over the strangers is the result of their having too much concern for their dog rather than too little concern for the strangers, their act of saving the strangers has moral worth despite the fact that their

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 at her other pertinent concerns as well. 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 in this situation was merely an accident, as it was merely fortuitous that she didn't have this other 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 their pertinent concerns.

What are the pertinent concerns? They are all and only those that will (or would) determine whether the agent 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. 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 accidentally right. For such a set of concerns 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 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

excessive concern for their dog would have led them to do the wrong thing in other situations. Now, I'll concede this point to Markovits provided that what interests us is whether the dog-lover is willing to sacrifice their 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 dog-lover's concern for their dog is excessive. And, in those contexts, their act would not count as having moral worth on the view that I'll be defending.

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, there isn't any simple way of spelling out what the relevantly similar situations are. Thus, 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 answers to such questions depend on the context in which they're being asked and on what's taken to be relevantly similar given our interests in those contexts. Thus, we could imagine one context in which we're interested in 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. But we could also imagine a context in which we're only interested in 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—remember we're assuming that, in this version of the case, she has a concern for all mammals, which includes the boy.

As I see it, this context-sensitivity is a feature rather than a bug. First, moral worth has to do with whether an agent was just lucky to have acted rightly, and, in general, whether someone counts as lucky is context sensitive. Take, for instance, the case of Kasamba. Like 90% of those living in his village, he contracted Ebola during a recent Congolese outbreak. Yet, unlike 75% of his fellow infected villagers, he survived. By comparison, far less than 1% of the world population ever contracts Ebola, and, of those that do, about 50% survive. Given all this,

we may wonder whether Kasamba counts as lucky. And it seems that there is just no unequivocal answer. In a context in which our interests make everyone in the world the relevant comparison class, we should think that Kasamba was extremely unlucky to have contracted Ebola and only somewhat lucky to have survived the infection. But, if instead our interests make his fellow villagers the relevant comparison class, we should think that he wasn't at all unlucky to have contracted Ebola but was very lucky to have survived the infection.

Second, this sort of context-sensitivity explains how it's possible for philosophers to draw opposite conclusions from the very same case: the case of Huck Finn's helping his friend Jim, a runaway slave, to elude authorities despite his believing that it is wrong for him to do so. Arpaly (2002) and Markovits (2010) conclude from this case that an act can have moral worth even if it wasn't motivated out of a concern to do what's right, for, as they see it, Huck's act was non-accidentally right. But Johnson King (2020) and Sliwa (2016) conclude the opposite: that an act cannot have moral worth unless it was motivated out a concern to do what's right, for, as they see it, Huck's act was merely accidentally right. Context-sensitivity explains this. In a context in which we're interested in whether Huck might be led to do wrong by his desire to help a friend and his willingness to do so even when he believes that it's wrong to do so, Huck's act counts as only accidentally right. For it seems that Huck just lucked out in that Jim was a fugitive slave rather than, say, a fugitive thief or murderer. But in a context in which we're interested only in whether Huck would ever be led by his appropriate concern for Jim's humanity to do wrong, his act counts as non-accidentally right, because this concern would never lead him to do wrong—at least, not when he's relevantly informed and it's combined with other appropriate concerns. So, whether we think that Huck's act was only accidentally right just depends on the context.

At this point, we should have a good enough grip on what moral worth is to be able to evaluate various proposed substantive views. I'll start by looking at two simple views in sections 2 and 3. 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. In section 4, I explain what this fundamental flaw is. And this leads me to introduce a new concept—the concept of *an ultimate moral concern*—in section 5. Lastly, in section 6, I employ this concept in developing a new account of moral worth and show how this account compares favorably to its rivals.

2. *The Simple Kantian View*

Most contemporary accounts of moral worth stem from the views of either Hume or 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” (G 4:396–401).

The Simple Kantian View: A right act has moral worth if and only if it was motivated out of a nonderivative desire/concern to do what's right (whatever that may be) and the belief that were they to perform this act they would do what's right.

The Simple Kantian View is subject to counterexample. 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.⁶ So, when the opportunity arises, he's moved to help the chimpanzee escape back into the wild out of a nonderivative concern to alleviate what he correctly perceives to be the chimpanzee's distress.⁷ 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. Yet, he ends up reporting the poachers to the authorities out of both a nonderivative concern for doing what's right (whatever that may be) and the belief that reporting lawbreakers to the authorities is his 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.⁸

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 nonderivative concern for doing what's right. Intuitively, though, it seems to be non-accidentally right, for it was motivated in a way that makes his acting rightly extremely reliable. After all, he had a concern not only to alleviate the chimpanzee's distress, but also to safeguard the welfare of both the

⁶ 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, 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 to which he's willing to assent), 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.

⁷ 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.

⁸ 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 instance, one might worry that Huck has too great a concern for being loyal to a friend. Consequently, Johnson King (2020, 195) worries that Huck could be led by such a concern to help Jim even if he were a fugitive murderer rather than a fugitive slave.

poachers and the authorities. He even had a nonderivative 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—at least, not if he were 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, of moral worth. So, contrary to the Simple Kantian View, we should deny that a right act's manifesting a nonderivative concern for doing what's right is necessary for that act to have moral worth.

Also, contrary to the Simple Kantian View, we should deny that a right act's manifesting a nonderivative concern for doing what's right is sufficient for that act to have moral worth. To illustrate, consider the following.⁹

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 care more about keeping their rather trivial promises to take their husbands golfing than about saving their patient's life. Consequently, each doctor is going to take her own husband golfing this afternoon regardless of what the other doctor is willing to 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, given the unwillingness of the

⁹ These are adapted from David Estlund's case entitled "Slice and Patch Go Golfing" (2017, 53).

other to participate in saving the patient, the best thing that she can do is to take her own husband golfing as promised. Now, each has a nonderivative concern for doing what's right (whatever that may be). But, unfortunately, this desire is not as strong as each's desire to keep her promise to her husband, 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 nonderivative concern for doing what's right. Consequently, they together produce what I've labelled in Table 1 as outcome O_4 —the outcome in which each goes golfing with her own husband while the patient dies painlessly.

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. Both Slice and Patch maximize 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 Slice nor Patch maximize 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 Slice nor Patch maximize 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. Both Slice and Patch maximize utility.

I'm assuming that, on the correct moral theory, one is required to keep one's promises whenever there's nothing better that one can do.¹⁰ And there is nothing better that either doctor can 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 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.

In *Golfing for Rightness' Sake*, each doctor was motivated to take her husband golfing out of a nonderivative desire to do the right thing (whatever that may be) and the belief that if she were to take her husband golfing she would be doing the right thing. So, on the Simple Kantian View, each doctor's act of taking her husband golfing has moral worth. But although each doctor did the right thing, neither of them deserve credit for acting rightly. For 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 is willing to do her part in saving the patient. So, we should reject the Simple Kantian View. Manifesting a nonderivative concern for doing what's right is insufficient to confer moral worth on an action.

3. *The Simple Humean View*

¹⁰ 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.

The other leading inspiration for accounts of moral worth is 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). Here’s a simple version of his view.

The Simple Humean View: A right act has moral worth if and only if it was motivated out of a nonderivative desire/concern to perform an act with right-making feature RMF (conceptualized as such) and the belief that were they to perform this act they would perform an act with RMF.

This view is also subject to counterexample, though what sort of counterexample 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’.¹¹ On this view, a right act will have moral worth if and only if it was motivated out of a nonderivative desire to maximize utility and the belief that were they to perform this act they 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 for the sake of argument 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. Given these assumptions, the non-fundamentalist version of the Simple

¹¹ Arpaly and Schroeder (2014) 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, 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, 87).

Humean View implies that my act of pushing the button would have moral worth if I were motivated out of a nonderivative 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 would have moral worth if I were instead motivated to perform it out of a nonderivative 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.

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 nonderivative concern to prevent the dog from getting hurt and that this is what makes her blocking the boy's blow the right thing to do. Of course, it's not what fundamentally makes it right. For assuming (merely for the sake of argument) that maximizing act-utilitarianism is correct, what fundamentally makes it right is that doing so would maximize utility.¹² 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. 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. Thus, the non-fundamentalist version of the Simple Humean View should be rejected.

The problem with this version of the Simple Humean View is that it fails to capture the counterfactual reliability that's required for moral worth. Consequently, it allows that the set of concerns that confers moral worth on your act is one that could lead you to act wrongly in other situations in which you are relevantly informed, making your acting rightly in this instance

¹² I'm assuming that the boy's blow would cause the dog a lot more harm if not blocked than it would cause her if blocked.

merely accidental. This is because the only way to ensure counterfactual reliability when relevantly informed is to consider not only the agent's motivating reason and whether it was good, but also whether the agent had all the other pertinent concerns and in the appropriate 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 nonderivative 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 (Manuscript) has pointed out, people rarely conceptualize their actions as meeting some fundamental moral criterion.¹³ In any case, *The Empathic* is a clear counterexample to the fundamentalist version of the Simple Humean View. Christoforos wasn't motivated out a nonderivative 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 nonderivative 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 also reject the fundamentalist version of the Simple Humean View.

¹³ See Star (Manuscript). See also Howard (Forthcoming).

4. Where These Two Simple Views Go Wrong

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 which version we're considering.¹⁴ 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 and the relevant knowledge have moral worth.

More specifically, the Simple Kantian View goes wrong in insisting that acts with moral worth must manifest a nonderivative concern for doing what's right, when, arguably, having a *nonderivative* 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 because we care about the ends at which morality ultimately aims and know that doing what's right is a means to our best furthering those ends. Thus, we should care about acting morally, not for its own sake, but only for the sake of the ends at which morality ultimately aims. Nathan Howard ([Forthcoming](#)) makes this 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

¹⁴ 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, 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*.

that we owe to our friends, family, and fellow humans.

Consider an analogy. Suppose I'm a sperm-donor who wants my progeny to flourish. Unfortunately, I don't know how many progenies I have, let alone who they are, for the sperm-bank insists on strict anonymity. Now, imagine that I can trust Episteme to tell me what I ought to do according to the means-end-rationality standard, and he tells me that I ought, according to this standard, to wire money into certain numbered accounts. What's more, I know that if I do what I ought, according to this standard, to do, my progeny will flourish. In this case, it seems that, although I should care nonderivatively about my progeny, I should care, only derivatively, about doing what I ought, according to this standard, to do. For doing what I ought, according to this standard, to do is merely a means to ensuring that my progeny flourish. This, I believe, is analogous to the situation in which I want to further the ends at which morality ultimately aims but am uncertain as to what those ends are. In such a case, I should care about doing what I morally ought to do only as a means to furthering the ends that I nonderivatively care about: the ends at which morality ultimately aims, whatever they may be.

The Simple Humean View also goes wrong; it does so in denying that an act can have moral worth in virtue of manifesting a derivative concern for doing what's right, when, arguably, such a concern is entirely appropriate. For suppose that I care nonderivatively about the ends at which morality ultimately aims and, so, want to do what will best further those ends. But suppose that I don't know how best to further those ends, either because I don't know what they are or don't know how best to further them. I may, nevertheless, know—as a result of, say, reliable testimony—what the right thing to do is. What's more, I may know that I can best further the ends at which morality ultimately aims by doing what's right. And, in such instances, I should have a derivative concern for doing what's right—one that's derivative of my nonderivative concern for the ends at which morality ultimately aims. Indeed, as 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, 408).¹⁵ Also, there will be times when one will be tempted to do wrong because one fails in the moment to be appropriately moved by the things that ultimately matter. And, in such cases, being moved to do what’s right can serve as a proxy for being moved directly by the things that ultimately matter.¹⁶ Yet, according to the Simple Humean View, an act has moral worth only if it manifests a nonderivative concern for its right-making features, and a derivative concern for an act’s being right isn’t the same as having a nonderivative concern for its right-making features. So, we must reject the Simple Humean View.

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 and the relevant knowledge have moral worth, where the appropriateness of a concern—both in terms of its magnitude and in terms of its being either derivative or nonderivative—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 nonderivative concern for doing what’s right), and the Simple Humean View fails in prohibiting us from having an appropriate concern (specifically, a derivative concern for doing what’s right). Of course, in making these arguments, I’ve relied heavily on the notion of what our ultimate moral concerns should be. Unfortunately, this notion has been undertheorized. So, I will need to take a brief digression from our discussion of moral worth to explicate it in the next section.

5. Our Ultimate Moral Concerns

Unfortunately, many contemporary moral theories fail to tell us what our ultimate moral concerns should be. That is, they fail to tell us what we, as moral agents, should ultimately be

¹⁵ Although she’s right about this, she’s wrong to assume that “what’s important is that the concern for doing what’s right be non-instrumental” (2016, 396). She claims that “the agent must care about doing what’s right for its own sake, and not because it would further some other goal” (2016, 396).

¹⁶ See Lillehammer (1997) for some nice examples.

aiming to achieve. There are at least three reasons why a complete moral theory owes us an account of this. First, the question of what we, as moral agents, should ultimately be aiming to achieve is itself an important moral question and, thus, one that a complete moral theory should 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, in turn, 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 achieve the ends that the theory directs them to achieve. And this is important, because we should reject such a theory.

To see why, 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. 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 they know that their 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. That is, it will require them to have incoherent motives. On the one hand, they'll be required to abide by (and to be motivated to abide by) the ideal code even when they know that doing so won't maximize the good. And, on the other hand, they'll 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, 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 Indeed, this is what he does. And 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 (and be motivated to 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. Now, admittedly, many contemporary moral philosophers have ignored this aspect of moral theorizing. They have often contented themselves with merely offering a criterion of rightness. Or if they go beyond that, they do so only to include a decision procedure. There are exceptions, of course. Nevertheless, contemporary moral philosophers have tended to neglect the issue of what our ultimate moral

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 requires ensuring that we act only in ways that are impartially defensible.

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 moral theories are potentially subject to this 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 aim 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.

We've seen, then, that the notion of an ultimate moral concern 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 determining whether the concerns that issue in an act are appropriate, which, in turn, is relevant to determining its moral worth. 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.

6. 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 and the knowledge (or, at least, the rational belief) that this is what would best further those concerns, where this set includes all and only pertinent concerns, each of which must be both qualitatively and quantitatively appropriate, which in turn depends on what the agent's ultimate moral concerns should be.¹⁸

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 depends on 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, they should be concerned with acting rightly only insofar as acting rightly is a means to furthering the ends for which they should have an ultimate moral concern. And, thus, their concern for acting rightly will be qualitatively appropriate only if it's derivative in this way. What's more, it will be quantitatively appropriate only if its magnitude is in direct proportion to the extent to which their acting rightly is, in the given situation, a means to their furthering the ends for which they should have an ultimate moral concern.

To take another example, imagine that we should have an ultimate moral concern for promoting each existing individual's utility but not for promoting the overall sum of utility in the universe. In that case, it would be inappropriate for us to want to bring happy individuals into existence for the sake of increasing the overall sum of utility, but appropriate for us 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, we should care just as much about n hedons of

¹⁸ I've added an epistemic condition here, because it seems that even if one has the appropriate set of concerns, one's act will count as merely accidentally right if one performs this act only because one irrationally (although, by luck, correctly) believes that this act is what will best further those concerns. I thank Ralph Wedgwood for pointing this out.

utility for one stranger as we do about n hedons of utility for any other stranger. Thus, how intensely we should want to promote the utility of an existing stranger by n hedons should depend solely on how great the number n is.

Now that we have a sense of how an agent's ultimate moral concerns determine the appropriateness of the concerns from which their 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. And there's a lesson here.

Lesson 1: A right act can lack moral worth even if it was performed for the right reasons—that is, for the reasons that make it right.

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 of 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).

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 and the relevant knowledge. In notifying the authorities and hiding the poacher's guns, Christoforos

manifests a concern (a derivative concern) for doing what's right as well as a concern (a nonderivative 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 nonderivative 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 nonderivative concern for doing what's right.

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

The Concerns View not only allows that an act can have moral worth without manifesting a *nonderivative* concern for doing what's right, but it also allows that an act can have moral worth in virtue of manifesting a *derivative* concern for doing what's right. And this is important, because, as many have pointed out, such a concern helps agents deal with both temptation and uncertainty. Sometimes, we're tempted to do what's wrong but are moved by our derivative desire to do what's right to resist that temptation. Other times, we do not know (or are uncertain about) what we should ultimately care about. But if we know what it's right to do and that by doing what's right we can best further the ends that we should ultimately care about, then we can do what's right as a means to furthering the ends that we should ultimately care about. So, although there is something problematic about being motivated out of a *nonderivative* concern to do what's right given that this is fetishistic (see Smith 1994), there's

nothing problematic about being motivated by a *derivative* concern to do what's right where one faces temptation or 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, a derivative concern) for doing what's right.

And this means that we should reject rivals of 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).

Lastly, cases such as *Golfing for Rightness' Sake* show us that being motivated to do what's right because it's right is insufficient to confer moral worth on an act. Consider that, in this case, 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 this brings us to our fourth and final lesson.

Lesson 4: A right act can lack moral worth even if its agent performs it because it's right.

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—views such as those defended by Johnson King (2020).

We've seen, then, that the Concerns View is superior to its rivals. 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.¹⁹

¹⁹ For helpful comments and discussions, I thank Ron Aboodi, Chrisoula Andreou, Cheshire Calhoun, Jamie Dreier, Josh Glasgow, Peter Graham, Brad Hooker, Nathan Howard, Paul Hurley, Jessica Isserow, Zoë Johnson King,

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.
- Baron, M. (1995). *Kantian Ethics Almost without Apology*. Ithaca, NY: Cornell University Press.
- 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.
- 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.

Andrew Khoury, Sarah McGrath, Aaron Rizzieri, Grant Rozeboom, Lucia Schwarz, David Shoemaker, Keshav Singh, Holly Smith, Hannah Tierney, Travis Timmerman, Mark Timmons, Ralph Wedgwood, and Erik Zhang.

- 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.]
- Lillehammer, H. (1997). "Smith on Moral Fetishism." *Analysis* 57: 187–95.
- Markovits, J. (2010). "Acting for the Right Reasons." *Philosophical Review* 119: 201–42.
- 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.
- Star, D. (Manuscript). "Moral Worth, Normative Ethics, and Moral Ignorance."