This is a draft. The final version is published in *The Journal of Ethics*:

https://link.springer.com/article/10.1007/s10892-025-09512-0

Please cite the final version.

The Moral Worth of Mixed Actions

Bowen Chan

Abstract: Ordinary people often act from both motives that are good and motives that are not. How should we assess the moral worth or value of these actions from mixed motives? Having neglected these actions, the recent literature leaves us with no obvious answer. In this paper, I develop an answer. A mixed action, I argue, can be morally worthy even if it is done neither purely from good motives nor partly from good motives that suffice in a relevant sense to prompt it. If and to what degree it can be morally worthy should be settled by weighing the goodness of its motives against their badness. And two properties of each motive, its intentional object and its motivational strength, together determine the degree of value that it contributes to the action's net worth.

Keywords: mixed motivation – moral evaluation – moral motivation – moral worth – W. D. Ross

An action's moral worth or value is commonly distinguished from its rightness. It is thought that two equally right actions can differ in their moral values by differing in their motives. Famously, Kant offers us the intuitive case of an honest shopkeeper, who gives honest change even to inexperienced shoppers solely as a means to avoid getting a bad reputation (G 4:393). His honest dealings, despite being right, have no moral value because his motive has none. They would, however, be morally good but no more right, if he had instead acted solely from, say, a concern for duty (or honesty) for its own sake. So an action is morally good, it seems, so long as

the motive that prompted it is morally good.[1] But we often act from multiple motives. Notably, we often act partly from motives that are good and partly from motives that are not. Now, how should we assess these actions from mixed motives, or what I call "mixed actions," where matters are not as clear-cut as in Kant's case?[2]

---

[1] For a longer discussion of an action's moral worth or value as distinct from its moral rightness (as well as the value of its consequences, see for instance Nomy Arpaly (2002, 69–71) and Thomas Hurka (2010). Although moral worth and rightness are commonly distinguished, most contemporary scholars think that morally worthy actions must be right. For a survey, see Zoë Johnson King (2018, 191–2). Some disagree. For instance, Julia Markovits says," [W]rong actions [can] have (at least some) moral worth" (2010, 240). Hurka too thinks that "whether [an action] is virtuous [or morally worthy] depends only on its motive" (2010, 63), which is a view that "was widely accepted in the late nineteenth and early twentieth centuries, for example by Hastings Rashdall, Franz Brentano, G. E. Moore, and W. D. Ross" (2010, 58). I too agree, but nothing in this paper hangs on this point.

[2] Although mixed actions are clearly possible and common, how to assess their value is under-explored. There is one paper on mixed motives in business ethics by Vincent Di Norcia and Joyce Tigner (2000). They argue that mixed motives are common in business, but in discussing how to evaluate their goodness, they note that goodness "is a fuzzy, gradated and complex notion, not a simple, distinct, binary concept," and leave the task of evaluation for "the situational judgement of decision makers and stakeholders" (2000, 8). Although goodness is, I agree, graded and complex, I will propose a more precise account that gives us clearer guidance for evaluating mixed motives from intuitive principles. Apart from that paper, Holly Smith (1991) has provided a great attempt at theorizing about an action's degree of moral worth in relation to the configurations of desires behind the action, and in this paper, I will consider and critique her proposed fitness model. Apart from these two papers, the recent debates about moral worth, however, focuses on asking which motives are good without seriously considering how to assess mixed actions.

Consider two cases of mixed actions.

Reputation: A nurse helps vaccinate patients during the worst of a serious pandemic, despite the risks of becoming ill, not simply because they care about promoting patients' health, but also partly because they want to be publicly recognized as a pandemic hero.

Bad Blood: An impatient patient tries to sneak in a meal right before an important blood test, but a nurse stops them. The nurse does so partly because they want to do their duty and partly because they want to get back at the patient for always asking too many questions.

It is hard to be certain whether each of their actions is good, and how good, if at all, each is. In each case, the nurse acts partly from a motive that is good and partly from one that is not. Judging their actions to be good simply because they act partly from a good motive is hasty. Judging their actions to be not good simply because they act from mixed motives is also hasty. We need more details.

Philosophers disagree about which motives are good and so may disagree about the cases. Some believe that only concern for rightness as such is good (e.g., Herman 1981; Sliwa 2016). So in *Reputation*, the action cannot be good for them because the nurse fails to act from duty. But in *Bad Blood,* the action can still be good for them because the nurse acts partly from duty. Others believe that only concerns for what makes actions (in a respect) right or good are good (e.g., Markovits 2010; Arpaly and Schroeder 2013). So in *Bad Blood*, the action cannot be good for them because the nurse fails to act from any concern for what makes actions right or good. But in *Reputation*, the action can still be good because the nurse acts partly out of concern for

promoting patients' health, which makes their action (at least in a respect) right and good. Whichever side is right here, how we should assess mixed actions remains an open question. In answering this, we can simply assume that whereas some motives are good, some are not. For the sake of neutrality, I refer to both types of motivating concerns above as good motives, But my arguments do not hang on this.[3] Feel free at any time to swap the types of motives.

Also, some philosophers may distinguish between motives and intentions or volitions and think that what matters in assessing actions are intentions or volitions rather than motives. Even so, they still need to explain how to assess a form of mixed actions, that is, those actions that are prompted by mixed intentions or volitions. In *Reputation*, the nurse can intend or will to both be publicly recognized and promote patients' health by vaccinating them. In *Bad Blood*, the nurse can intend both to do their duty and to get back at a patient by stopping them. And we can again ask how should we assess a mixed action. But for simplicity's sake, I assume that what matters in assessing an action are simply motives, and in particular effective ones, or those motives, which may be further understood as either motivating reasons or desires, that in fact help cause

[3] There is also an emerging trend of pluralism that counts both types of motivating concerns as good (e.g. Hurka 2014b; Johnson King 2019; Tomlinson 2019; Isserow 2020). And as an interesting historical point, the pluralist view of good motives was popular and widely held by British ethicists, such as W. D. Ross, during the early twentieth century (Hurka 2014a). Now, even if one holds both types of motives are good, people debate about whether good action requires acting at least partly from concern for rightness as such: some argue it is always needed (Johnson King 2018), some argue it is sometimes needed (Aboodi 2016), and some argue it is never needed (Howard 2019). And of course, depending on which view here is right, the scope of good mixed actions will vary, but my analysis will be compatible with either of these views.

or explain the action.[4] So I refer exclusively to effective motives.

The question of how to assess mixed actions intersects with many other independent and orthogonal controversies about which I remain neutral. In particular, I am also neutral about whether only right actions can be morally good as well as relatedly whether only actions done from the belief or even knowledge that one is acting rightly are morally good. And by "right," I am neutral about whether it should be understood in purely consequentialist terms. Whichever answers to these questions are right, my analysis of mixed actions is still relevant.

I begin by considering two possible models for assessing mixed actions. According to one, "[t]he *purity* model" (Smith 1991, 281), a good action must be untainted by what is not good, so must be prompted purely by good motives. Given this model, mixed actions are never good. According to a second, the "fitness" model (Smith 1991, 281–2), a good action must be prompted by good motives that are sufficiently fit. More specifically, a mixed action is good if and only if it is prompted at least partly by a good motive that suffices in a sense to prompt it. (When we get to examining this model, we will consider three different readings of "suffices.") While initially appealing, these models, I argue, have unacceptable consequences.

I defend a neglected alternative model. Just as when assessing consequences on the whole, the badness of one bad consequence may not outweigh the goodness of all the good ones, when assessing an action on the whole, one bad motive may not outweigh all the good ones. Bad motives always make actions in a respect bad. But their badness can be outweighed by good

[4] Perhaps, in assessing actions, we should also count motives that are present but that one does not act on or ones that are absent but that one should have acted on. It will help to consider whether they do count in providing a fuller account of moral worth, but for this paper, I will focus on mixed effective motives.

motives that also always make actions in a respect good. And just as some consequences are neither good nor bad, and so make the overall state of affairs neither better nor worse, there are neutral motives that are neither good nor bad, and that (at least by themselves) make an action neither better nor worse. For an overall assessment of an action, I argue that we should weigh its total goodness against its total badness. Call this "the weighing model." Assessing whether, and if so, to what degree, an action is good on the whole, I argue, requires weighing its entire set of *effective* motives, that is, those motives that actually helped prompt it. A sketch of this model is offered in the last and much-neglected pages of W. D. Ross's (1930) *The Right and the Good*, which deserves some serious attention. Ultimately, the weighing model, I argue, captures our intuitions better than others. Moreover, this model is not only interesting in its own right in illuminating the value of mixed actions, but also shows how assumptions about the nature of motivation bear on whether neutral motives, which initially appear to make an action neither better nor worse, can result in an action that is less than ideally good.

1 The Purity Model

Deciding when mixed actions such as in *Reputation* or *Bad Blood* are good can be difficult. This is so because they are underspecified. But once we further specify them, it will be clearer that, contrary to the purity model, mixed actions can be good.

First, consider again *Reputation*. This is a case where one acts both from a good motive and from what I call a "neutral motive." Besides good motives, such as a non-derivative motive to please you or to do what is right, and bad motives, such as a non-derivative motive to hurt you or to wrong you, there are neutral motives. A non-derivative motive to get others' approval, for example, is neutral according to the two dominant views about which motives are good (and

what I take to be their implicit views about which motives are bad). For those who think that only concern for rightness as such is good, a motive to get others' approval is not good. Conversely, it is not bad if only concern for wrongness as such is bad. Likewise, for those who think that only concern for features that make actions right is good, it is not good. Conversely, it cannot be bad if only concern for features that make actions wrong is bad. Now, it is clear that a motive to get others' approval for its own sake is simply neutral, neither good nor bad. So since in *Reputation* the nurse is not acting from a bad motive, their action is at least not bad, and since they act partly from a good motive, their action is at least in a respect good. Although they are not acting purely from a good motive but also partly from a neutral one, their action, it seems, is still on the whole good.

Amy Mullin (2011, 2016) makes a similar point in discussing gratitude to paid caregivers. She also observes that "we typically have mixed motives for much that we do" (2016, 726). And more specifically about paid caregivers, she adds,

> [I]f recipients have reason to believe that care labor directed at them is intended not only to make ends meet for the laborer but also that the care laborer has also chosen a field of work in which he (or more likely she) can meet real needs in a respectful manner, then the paid laborers merit gratitude. (2011, 118)

In benefiting you, an agent deserves your gratitude if they did so partly because of benevolence or some other good motive (at least as long as it is also done so respectfully). Acting also partly from neutral motives, like a motive to grow their bank balance, however, does not diminish at all (at least not by itself) the goodness of their otherwise good action. This point helps explain why

often we should still be grateful to our paid caregivers.

Of course, you may worry that if they act mostly from a very strong neutral motive and only somewhat from a weaker good motive, the neutrality of the neutral motive will eclipse the goodness of their otherwise good action, thus making it on the whole without any value. For instance, if a nurse acts *mostly* from a very strong concern for getting public recognition and only *somewhat* from a concern for the patients, their action, it seems, would not be good. An agent may act from the same motives, but the strengths of each of their motives may differ, A year ago, the agent may have cared a lot more about themselves than about others, and so acted mostly from self-interest than from concern for others even when helping others then. But today, the agent may care as much about themselves as about others, and so now act equally from a concern for themselves as for others when helping others. We can understand the strength of each motive or concern roughly to be the range of counterfactual situations in which given the actual motive with its actual strength, the motive would in fact be effective. Whether any motive would still be effective depends on a constellation of many other factors including whether their cooperating and countervailing motives stay fixed as they in fact are. So it is quite difficult to determine the exact range of situations in which it remains effective. But for our purposes, we can do with just a comparative measure of its strength. For example, if I say a nurse acts partly from a strong neutral motive, say, with a strength of 50 to help, and partly from a much weaker good motive, say, with a strength of 5 to help, I am roughly saying that they would still help if they had acted solely from the neutral one in many more situations than they would still help if they had acted solely from the good one. And more specifically, the nurse would have still helped if they had acted solely from the good one only in situations where their countervailing motive had a strength ranging from 1 to 4 whereas in comparison they would have still helped if they had

acted solely from the neutral motive even in situations where their countervailing motive had a strength ranging from 5 to 49. Considering this case, the nurse's action is not so obviously good because their neutral motive is so much stronger.

Those who think that intentions or volitions are all that matters in assessing actions may dismiss the worry because for them a motive's strength is irrelevant. But it seems to me that intentions and volitions also differ in strength. We do speak of strength of will and willpower. And as I have noted before, intentions and volitions can be mixed. When they are also mixed, we can ask how strong are each of your intentions or volitions. And again, we can understand the strength of each intention or volition roughly to be the range of counterfactual situations in which given the actual intention or volition, you still follow through with your action. Actual success may not even matter, because actual trying is enough. But we can still assess your trying by considering your intentions or volitions and their respective strengths. Now, we may again worry that an agent's neutral intentions or volitions can eclipse their good ones, thereby making their action without any value.

Although neutral concerns may eclipse good ones when they are much stronger than them, and the good concerns can also equally eclipse the neutral ones. Suppose in a modified case, *Reputation+*, a nurse keeps vaccinating patients despite the personal risks, *somewhat* because they care about getting recognition, but *mostly* because they care about helping patients. Now, since in *Reputation+* the neutral motive is neutral, and the good motive is not only good but also the stronger motive, this nurse's mixed action, I think, is clearly good on the whole.

Second, even acting from good motives mixed with bad ones is at times good on the whole. Suppose in a modified case, *Bad Blood+*, a nurse stops a patient from eating before a blood test *mostly* for the sake of doing their duty but also *a bit* for the sake of getting back at the patient.

They act partly from a bad motive, so their action is in a respect bad. But it is not wholly bad. They also act partly from a good one, so their action is in a respect good. More importantly, they act *mostly* from a strong good motive and *only a bit* from a much weaker bad motive. Hence I think that their action is still good on the whole.

Of course, I am not saying mixed actions that are good are always good to the same degree. The good actions of every nurse (or person) are definitely not all good to the same degree. Some nurses act from a stronger concern for their patients or their duty to them than others, and other things equal, the actions of the former are definitely better than those of the latter. Some do on occasion act partly from a bad motive, say, petty spite, while others do not, and other things equal, the actions of the former are definitely less good than those of the latter. Some may act from a stronger neutral motive, say, concern for a form of reward, than others, but if their other motives were entirely the same, the actions of each may be equally as good.

According to the purity model, however, none of these mixed actions can ever be good, According to the model, an action is ever good only if it is done purely from good motives. Barbara Herman presents a Kantian version of the model:

> [W]hen an action has moral worth, nonmoral *incentives* may be present, but they may not be the agent's motives in acting. If the agent acts from the motive of duty [i.e., a good motive], he acts because he takes the fact that the action is morally required to be the ground of choice. It does not follow from this that the action's moral worth is compromised by the presence of nonmoral feelings or interests, so long as they are not taken by the agent as grounds of choice: as motives. … Strictly speaking, the doctrine of moral worth can accept the overdetermination of action with respect to *incentives*, not motives. (1993, 12)

On her view, one can have multiple incentives (motives that are not but could be effective), but for an action to be good, only good motives can be effective, or help prompt the action.[5] Neither bad nor neutral motives can be to any extent effective, that is, help prompt the action. So mixed actions are never good

But as our earlier discussion of nurse cases suggests, at least some mixed actions are good. More generally, if you act solely from good motives, your action, it seems, should be good. Even if you also act partly from neutral ones, your otherwise good action should still be good. But if you also act from bad motives, your otherwise good action would be in a respect bad, though not necessarily bad on the whole as the original goodness can outweigh the badness.

More importantly, some differences in the degrees of value can exist between mixed actions that are done simply from a mix of good and neutral motives, untainted by any bad ones. Many of these actions at least are clearly not bad. Now, if mixed actions can never be good, there cannot be any difference in the degrees of value between these cases of mixed action; they must simply be without value. But as we saw earlier, there can be a difference in value. Compare the nurse who acts more for the sake of reputation than for the sake of their patients with the nurse who acts more for the sake of their patients than for the sake of reputation. Clearly, the former's action has less value than the latter's. So a difference in value can exist. Now, since the purity

[5] Herman's use of "incentives" just refers to "nonmoral motives" that are "present and yet not operative [i.e., not what moves the agent to act]" (1993, 11). And of course, Herman believes, among other things, that the motive of duty is the only good motive. And many followers of Kant also likely believe this idea and the purity model (e.g., Stratton-Lake 2005). But some may believe the purity model without thinking that the motive of duty is the only good motive or that it is a good motive at all.

model cannot account for the intuitive possibility of such a difference, we should reject the purity model, or at least the version we have been discussing so far.[6]

A weaker version of the model is correct, however, given one possible view of motivation. We have seen that an action's goodness depends on the strength of its effective good motives and its badness on the strength of its effective bad motives. We may think that its net value depends more specifically on the strengths of those good and bad motives. And there are, however, two possible views of (effective) motivation. On what I label as the "capped view," the effective strength of one's motives taken together is equal to the bare minimum needed to overcome the total strength of countervailing motives, which may vary across situations.[7] Again, suppose a nurse is asked to come to work to help vaccinate patients during a pandemic. For ease of illustration, suppose they need a motive with a strength of 5 to help do so, but because they have a stronger motive with a strength of 7 to help, they do come in to help. According to the capped view, the effective strength of their motive to help is simply 5, which is the minimum needed to

[6] Wim Dubbink and Luc van Liedekerke claim "moral purism" is actually "a common way of thinking [that is, a pre-theoretical cultural view] about the moral quality of business people's actions," even though it is rare in academic work as "it is not (fully) explicated discursively" (2019, 379). But it is unclear what people's purism amounts to even if they are cynical about business people. People's cynicism here may just reflect implicit bias towards business people. In particular, they may tend to assume business people, and profit-driven people, to have bad ulterior motives, without any genuine good motives at all, or at least without any strong ones. If true, the typical cynicism here may just reflect beliefs about their motives rather than the belief that only actions from purely good motives are ever good.

[7] In principle, it may be also be capped by a constant ceiling across all situations. But it is likely false since overcoming certain temptations, say to steal, are often easier than others, say to procrastinate.

overcome a countervailing motive not to help.

This view will appeal to those with a certain view of explanations or causes more generally. Suppose a tower that would collapse if it is hit by an earthquake of at least magnitude 5 gets hit by a massive earthquake of magnitude 7. Some may think that it collapsed simply because it was hit by an earthquake of at least magnitude 5; the fact that the earthquake was stronger than 5 by 2 specifically contributes nothing to the explanation or cause of the collapse. Similarly, it may be thought that the total effective strength of one's motivation is capped by the minimum needed to overcome one's countervailing motives.

Given the capped view, neutral motives do indirectly make an action less good overall. They take up a limited supply of effectiveness that could be instead given to good motives.[8] Suppose A acts from a mix of just good and neutral motives against some countervailing ones. As A's neutral motives become more effective, A's good ones must become less effective, and A's action thereby becomes less good. Conversely, as A's neutral motives become less effective, A's good ones must become more effective, and A's action thereby becomes more good. Neutral motives still cannot make actions bad on the whole, only less than ideally good. Now, given the capped view, a weaker version of the purity model can be correct in claiming that the best actions are done from purely good motives.[9] Still, this weaker version must accept that mixed

[8] This claim also assumes that enough good motives are always available to become effective motives by taking up the limited supply of effectiveness (and that bad motives are not the ones taking up that supply).

[9] W. D. Ross has suggested that Kant thinks that "the worth of an action is degraded by the presence in it of any motive lower than the sense of duty [i.e., the best motive]" (1930, 170), because "he is assuming that the motivation is always exactly enough to produce the doing of the given act" (1930, 172).

actions can be good without being ideally so.

But on the uncapped view, the effective strength of one's motives is *not* capped and includes both the minimum needed to overcome countervailing motives and any surplus motivation, beyond that minimum that is present in one's effective motives. For illustration, consider again the nurse being called to work to help during a pandemic. According to the uncapped view, even though they just needed a motive with a strength of 5 to help, the effective strength of their motive to help is 7, which includes not just the minimum needed but also the surplus.

This view will appeal to those who think that the tower collapsed just because it was hit by an earthquake of magnitude 7. One of a mere magnitude 5 would have also caused a collapse, but the earthquake of magnitude 7 in fact caused it. This parallels the view that there is no cap on the effective strength of one's motivation, where the effective strength includes the surplus beyond what is needed to overcome countervailing motives.

Given the uncapped view, neutral and good motives need not compete for effectiveness. With no cap on the total effectiveness, an increase in the effectiveness of one's neutral motive need not result in a decrease in the effectiveness of one's good one. The two are independent. Suppose B's neutral motive is much stronger than C's. From this alone, we cannot assume that B's good motive's effective strength is weaker, and B's action is thereby less good, than C's. Given an uncapped model, B's good motive could be effectively as strong as or even stronger. So B's action need not be less good simply because it is done from stronger neutral motives.[10] So purity need not be better.

---

10 Ross has suggested this view. He says that in acting from both duty and a neutral motive, there may be a "surplus of motivation," where "duty is strong enough to have secured by itself the doing of the act, and that in that case the action is as good as if it had been done from sense of duty alone" (1930, 172)

Whichever view of motivation is right, however, mixed actions can be good to some degree, contrary to the purity model's central claim that they are never good, which is simply wrong. Of course, if the capped view is right, the best actions are done purely from good motives, and the purity's model weaker claim that purity is always better will be right.

2 The Fitness Model

Another model claims that your action is good if and only if you act at least partly from good motives that suffice to overcome countervailing motives, if any, and prompt you to act. Holly Smith has called this the "fitness" model (1991, 281–2). There are three versions of this given three ways of reading "suffice" (1991, 284–90).[11] On one reading, which Smith prefers, you simply need act from a good motive that is part of a set of motives that prompt you to act. Call this the "minimal condition." Your effective good motive need be neither *sufficient* nor *necessary* for prompting you to act. On a second reading, they need to be *sufficien*t for action, and thereby stronger than your countervailing motives. Call this the "sufficiency condition."[12] As on the minimal condition, your other motives are irrelevant, but unlike on that condition, your effective good motives must be stronger than countervailing motives. On a third reading, your effective good motives must be *necessary* for action, that is, without their cooperation, your

[11] Smith also gives a fourth and final version: both sufficiency and necessity conditions must be satisfied. I set this aside because I argue that the sufficiency and necessary conditions are each unnecessary, and that suffices to show that their conjunction is unnecessary. Also, Smith herself does not endorse it, mentioning it only for completeness's sake.

[12] In her paper, Smith calls this "strong sufficiency" and the next one "weak sufficiency" (1991, 285), but calling them "sufficiency" and "necessity" should be clearer since their key difference is not in degrees.

other effective motives cannot be sufficient for action. Call this the "necessity condition." Unlike the first two, it also imposes a restriction on your neutral and bad effective motives: these motives must be weaker than countervailing motives.

Unlike the purity model, the fitness model allows some mixed actions to be good so long as their effective good motives in some sense suffice. But I will argue that it is still inadequate because each of its versions has counterintuitive consequences.

2.1 Minimal Condition

I agree that satisfying the minimal condition is necessary for an action to be good because without a good motive, an action clearly cannot be any good. But satisfying it is not sufficient. Acting just partly from a good motive is not always enough to make an action good overall.

I may act rightly and do so *partly* from a good motive, but if I do so *mostly* from a bad one, my action can still be bad overall. Consider a case.

> *Call of Duty*: An officer knows that they must use force to arrest a clearly violent criminal. The officer genuinely cares about doing their duty to keep their neighbourhood safe and they also enjoy the use of force, and they do so much more than they care about their duty. Despite this, they are also really afraid of getting hurt, and so given either motive alone, they would have been too terrified to confront the criminal. But because they want to both do their duty and use force, they still arrest the criminal.

Suppose the officer cares more about doing their duty than the average civilian does, but without also enjoying violence, they would have run away. Their action would be clearly bad if they

acted solely from a strong bad motive. Although they also act partly from a good one, their much weaker good motive fails to fully redeem their otherwise bad action. Of course, even if several officers act partly from a bad motive, their actions need not all be equally bad. One officer may act from a stronger desire for violence, and so act even worse, than a second, but both their actions may be less bad than a third who acts solely from a desire for violence.

Whereas the purity model is too demanding in denying moral value to any mixed action, the minimal condition is too permissive in granting it to too many such actions. On one hand, good motives make actions in a respect good, and stronger ones can make them even better. On the other, bad ones make actions bad in a respect and can even make them bad on balance. So acting partly from a good motive is not sufficient for an overall good action.

2.2. Sufficiency Condition

To accommodate *Call of Duty*, the fitness model can be read with the sufficiency condition according to which good actions cannot depend at all on bad or neutral motives to be done: your action is good if and only if your good motives *sufficed by themselves* to bring it about.[13] When your good motives do not suffice to, or could not by themselves, prompt you to act, your action will have to be done partly from other motives that may be bad (or neutral).

Depending on bad motives to act, however, need not make your action bad on the whole. Imagine a variant of *Call of Duty*. The officer cares a great deal more about their community and enjoys violence a great deal less. So they stop the criminal mostly from a sense of duty and only minimally from enjoying violence, though neither motive on their own is sufficient. Their action, I think, can still be good on the whole. They need the extra incentive of violence, but they need it

[13] Kant has also been read as holding this condition (e.g., Henson 1979).

only because they rightly fear getting hurt. Compare them to a new officer who does not need it only because they are barely concerned with their safety, and as a result, their sense of duty alone suffices. But the new officer cares just as much about their duty and enjoys violence just as much as the first officer. The new one also acts from mixed motives – from a sense of duty and an enjoyment in violence. Their action seems equally as good as (or comparable to) the first's insofar as the new one's effective motives are the same as the first's. It is unfair to judge the new one's action as good, but the first one's as entirely without value just because the new one is less scared. The first is more scared but is still brave insofar as they too act mostly from a strong sense of duty, and only minimally from enjoying violence.

Also, whether their good motive by itself suffices for action still depends on another motive, namely, their countervailing motive to avoid getting hurt given their fear, and how strong it is. The first officer's good motive is insufficient in a respect since they would not do their duty if they enjoyed violence less, but the new officer's good motive is insufficient in another respect since they would also not do their duty if they were more fearful of getting hurt. So we can see that whether their good motives suffice for action depends not only on a cooperating motive, their attraction to violence, but also on a countervailing one, their aversion to getting hurt. Now, it is arbitrary to say that an action is good even if it depends on countervailing motives, but it cannot be good if it depends on cooperating motives. Both depend similarly on luck. One may say instead that a good action must not depend on either of these kinds of motives. But if this is true, no action can be good. No good motive shields one from all temptations.[14] Instead, we

[14] Even if what matters is not a good motive but a good intention or volition that is wholehearted, wholeheartedness also depends on not being affected by any temptation, which again depends on luck.

should accept that independence from bad motives is not needed for good action. An action can be good even when it depends on bad motives to be done.

Moreover, even when good motives are insufficient, an action may depend not on bad ones but simply on neutral ones, which do not make an otherwise good action entirely valueless. Imagine a second variant of *Call of Duty*: the officer is partly prompted by their sense of duty but also needs the added incentive of getting recognition rather than of enjoying violence. Also, we can imagine their good motive is stronger than their neutral one (but we need not). Here, their action can still be good on the whole despite depending on a neutral motive.

Consider individuals who do work that involves caring. When nurses care for patients, their actions are good if they do so with real concern for them, despite also depending on neutral incentives, such as money or recognition, as others in the service industry do. Restaurant servers may care about customers but may not serve them from kindness alone. They may also need the incentive of tips. When they serve customers partly from real care, their actions are at least to some degree good, despite also depending on a neutral incentive.

Some deny that good actions can depend to any degree on neutral motives such as money. Good actions can be done partly from neutral motives, but they cannot depend at all on them. But this is, I think, a counterintuitive standard to hold people to, when in our non-ideal world, many people must spend most of their lives making a living wage to thrive and even survive. So if service jobs had no monetary incentive, many of these jobs would require a big sacrifice. But intuitively, making a big sacrifice is not necessary for good action.

For a related discussion of counterfactuals and moral worth, see Smith (1991), Steven Sverdlik (2001), Philip Stratton-Lake (2005), Kelly Sorensen (2014), Paulina Sliwa (2016), Jessica Isserow (2018), Johnson King (2018) and Nathan Howard (2019).

Whether or not you agree with me that this is a counterintuitive standard to hold others to, you should at least accept that actions from good and neutral motives can differ in their value, and once you do, you should, I argue, have strong reason to reject the sufficiency condition. Suppose B and C are officers who each stop a violent criminal despite fears of getting hurt. Together their concerns for duty and for getting good reviews are just enough for them to act, and they are both equally scared of being hurt. So their total effective motivation is the same. Also, suppose B depends more on the prospect of getting good reviews to act than C does. Now, since they have the same total motivation, B's good motive must be weaker than C's, and B's action must therefore be less valuable than C's. So clearly in this case, the action that depends more on neutral motives, B's, is less valuable than one that depends on them less, C's. So wherever they are on the scale, it is implausible to think that they rest on the same spot.

But the sufficiency condition cannot accommodate differences between B's and C's actions. According to the sufficiency condition, their actions are mixed and so cannot be any good. Now, if they cannot be any good, they must be either positively bad or completely valueless. Consider the first option. If they are both bad, they can differ as B's can be worse than C's. But it is quite implausible to insist that officers who need neutral motives must act badly, especially if their neutral motives are very weak, and their good motives are very strong.

A second option is to insist that B's and C's actions have no positive or negative value at all. All actions done from good motives with the needed help of neutral ones just have zero value. But if so, we cannot accommodate the above comparison between B's and C's mixed actions as one valueless action can be neither more nor less valuable than another equally with no value. The sufficiency condition is less restrictive than the purity model, but it agrees with it here and for similar reasons fails to account for differences in the degrees of value in mixed actions that

are done solely from good and neutral motives.

So the sufficiency condition cannot account for the differences in B's and C's mixed actions unless we accept the implausible claim that needing a neutral motive must make an action bad. We should therefore reject the sufficiency condition. Instead, we should embrace the claim that an action can be good even when it depends on motives other than good ones to be done.

We should not only deny that satisfying the sufficiency condition is needed for good action but also that it suffices for it. Even when good motives by themselves suffice to prompt action, an action can still be bad on the whole. Consider a third variant of the original *Call of Duty*. Suppose the officer's good and bad motives each alone sufficed to prompt them to do their job. They would still have stopped the criminal even if doing so did not require any violence, and they would still have stopped that very person even if they did not suspect them of any crime. If they acted partly from a concern for duty but mostly from a stronger appetite for violence, their action would seem bad on the whole.

So independence from bad motives matters insofar as depending on bad motives is bad. But complete independence, I have shown, is neither necessary nor sufficient for good actions. Actions can depend on bad motives but still be good overall if done mostly from good ones, and they may not depend on bad motives but still be bad if done mostly from bad motives. And independence of neutral motives is even less important. It essentially matters only when your dependence on neutral motives is excessive relative to your dependence on good ones. The sufficiency condition is therefore too demanding and too permissive. It is too demanding as it requires that your effective good motives be self-sufficient, finding no value in actions that depend on neutral or bad motives even when its good motives are overwhelmingly good. It is also too permissive as it only requires that your effective good motives be self-sufficient. thereby

finding goodness in an action even when its bad motives are overwhelmingly bad.

2.3 Necessity Condition

An alternative to both the minimal and sufficiency conditions is the necessity condition, which says your action is good if and only if good motives were *needed* to prompt your action, so your action could not be done without their cooperation. It denies goodness to cases where one's bad effective motives alone suffice for action such as the third variant of *Call of Duty,* where although the good motive alone suffices, the stronger bad one alone also suffices. Also, it grants goodness to cases where one's good effective motives alone suffice for action such as the first and second variants of *Call of Duty*, where despite needing a bad or neutral motive, one still acts partly from an overwhelmingly good motive.

But the necessity condition wrongly denies goodness to cases where one's good motives are not necessary but are self-sufficient for action. A mixed action can still be good in these cases. Imagine a fourth and final variant of *Call of Duty*: the officer stops the criminal both because they want to do their duty and to get good reviews, and either motive on their own suffices, but suppose their concern for duty is also more effective than their concern for good reviews. Despite not depending on the good motive to act, their action, it seems to me, is still good.[15] The good motive is really good, and also the more dominant motive, even if it is not needed. And the

[15] Even if only one of the two motives could have prompted their actual action, it still remains false that, one's good motives must have been needed for one to have actually acted for one's action to be good. Suppose their concern for the duty to keep the neighbourhood safe by itself prompts them to do their job. Their action still seems good despite their good motive not being needed for them to have acted.

neutral motive is not bad and so should not make the otherwise good action valueless. So dependence on good motives for action is not necessary for good action.

And the necessity condition wrongly grants goodness to cases where one's good motives are necessary but also depend on overwhelmingly bad motives. Recall the original *Call of Duty*. The officer needed both good and bad motives to act, but their action was still on balance bad because the bad motive, being much stronger, was much worse than the good one was good. Although the fact that the good motives were needed guarantees that they were also effective, and so even when one also acts from bad motives, it is a sign that one has redeeming qualities, it does not guarantee that overwhelmingly bad motives were not also needed and effective. Dependence on good motives for action is therefore also not sufficient for good action.

So although the necessary condition accommodates the first three variants of *Call of Duty* which the sufficiency condition fails to do, it fails to accommodate the original *Call of Duty* and also a new fourth variant, which the sufficiency condition can successfully accommodate. Thus each version of the fitness model fails to capture the full spectrum of mixed actions.

3 The Weighing Model

Having surveyed a spectrum of actions, here is my proposal: an action is on the whole good only if the degree to which the good effective motives are good is greater than the degree to which the bad effective motives are bad.[16] We may act from good, bad, and neutral motives.

[16] Again, some may believe that it is not first-order motives but higher-order volitions or intentions that count in assessing actions (e.g., Sverdlik 2001). But even these higher-order attitudes can be mixed, and an action is still good only if the goodness of the good attitudes outweighs the badness of the bad ones.

Good motives make an action *pro tanto* good. Bad motives make an action *pro tanto* bad. Neutral motives at least by themselves make an action neither good nor bad. An action is good only if and to the degree its *pro tanto* goodness outweighs its *pro tanto* badness. Call this model "the weighing model."[17]

The model does, however, ignore possible contributions to an action's value from countervailing motives. There are at least three plausible views about their contributions: (1) they contribute no value to actions, (2) facing temptations makes actions better, or (3) facing temptations makes actions worse. For this paper, I am focusing on mixed actions, or actions done from a mixture of motives that help prompt it, and so remain neutral on the question of how temptations affect an action's value and the corresponding debate between the defenders of the battle-citation model and the harmony thesis. To read more on this question, see for instance Richard Henson (1979), John McDowell (1979), Smith (1991), and Karen Stohr (2003). Even though the model fails to address this question, the model can be supplemented by answers to it. More importantly, whichever view is right, the weighing model's proposal that an action is good only if the goodness of its effective good motives outweighs the badness of its bad ones should still be accepted because the motives that you in fact are acting from matter more than the ones that you are overcoming. Even if you overcome temptation, it should not make up for acting from motives that are on balance bad

[17] Laura Tomlinson (2019) has also gestured in passing at the idea of weighing. In discussing Huck Finn, who helped his friend Jim escape slavery under the false impression that what he was doing was wrong, she suggests that "Huck's action is a case of *helping a friend*, *keeping a promise*, and *doing something he believes is wrong*. The goodness of the first two [action-]types clearly outweighs the badness of the third" (2019, 713). Though we share the idea that we can weigh good and bad parts of actions, I should note that I disagree, however, with the analysis of Huck here and what she is counting as a bad part of his action. Huck's action is not truly a case of acting from good and bad motives. His internal struggle suggests that

This draws on Thomas Hurka's (2001) preliminary model for assessing the virtuousness or moral value of a combination of attitudes, which says "the value of a combination of states always equals the sum of its components' values" (59). Just as Hurka's preliminary model takes the value of a whole to be the sum of its parts, the weighing model also does the same.[18] But whereas Hurka's model is concerned with assessing all attitudes present in one's mind, and not only those that are effective in action, the weighing model is concerned exclusively with those that are effective. According to the weighing model, when a nurse helps a patient, only effective motives, say, concern for the patient's health and their own reputation, count, whereas other attitudes, say, a longing for cheap rent or a hatred of squirrels, are irrelevant. Also, according to this model, when a judge rightly sentences a criminal only out of malice, despite having a strong

---

he does not want to do the wrong thing, which is a good motive. When he does what he thinks is wrong, he does not actually act from wanting to do what is wrong; he does it because he wants to help his friend, and because he wants to keep a promise. A conflicted Huck acts better than one who is not worried at all by the fact that, to him, he is acting wrongly. Hurka (2014, 498-9) makes the same point about Huck.

[18] Hurka (2001) rightly argues, however, that this preliminary model fails to account for proportionality and supplements it with a proportionality principle. The weighing model, I think, should also add one. Even so, the weighing model as described so far is still accurate as a necessary condition for good action. If the bad motives outweigh the good motives, one's relevant set of motives is definitely disproportionate, and one's action is even more clearly bad on the whole. Of course, the model as described so far misses the extra respect in which a mixed action will be bad insofar as its motives will be disproportionate. Still, it captures well the other respects in which it will be good or bad in virtue of its effective motives.

but in this case ineffective sense of duty, only the malice counts.[19] Thus the weighing model modifies Hurka's model to assess mixed actions. In these respects, the weighing model matches one model of acting from mixed motives that is sketched by Ross in the final few pages of *The Right and the Good* (1930), which I will discuss in the last section.

3.1 Strengths and Objects

A simple version of the weighing model may say an action is good only if and to the degree the good effective motives are effectively stronger than the bad effective motives. This is true if a motive's degree of goodness or badness is determined just by its strength. But this is false. Besides its strength, a motive's degree of value is partly determined by its intentional object, or what is being pursued for its own sake. A weak bad motive can be worse than a stronger one. A desire to (wrongly) kill is worse than an equally strong desire to (wrongly) worsen a cut. And the former bad motive is worse than the latter even when the former is slightly weaker. So a bad motive's degree of badness depends not simply on its strength but also on its object. A similar point fits with good motives. A weak good motive can be better than a stronger one. A desire to (rightly) save a life is better than an equally strong desire to (rightly) treat a cut. And the former

[19] Ineffective motives can, I think, make actions less good and even at times on balance bad. For instance, If A, say, hires B just out of love for B, with a much weaker ineffective concern for being fair to C and D, A acts badly despite acting on a good motive (Ross 1939, 307). But it will exceed the scope of this paper to explore how ineffective motives can affect an action's value. And even if ineffective motives can count, it remains true that an action is good only if the total goodness of the good effective motives outweighs the total badness of the bad ones. Rightly punishing a criminal solely from malice is still on the whole bad even if one has a very strong ineffective sense of duty.

good motive is better than the latter even when the former is slightly weaker. So like bad motives, a good motive's degree of goodness also depends on its object.[20]

For an action to be good, good motives need not be effectively stronger than the bad ones, but a given motive with the right object cannot be good if it is too weak in absolute terms. Suppose a person cares only a bit about giving to charity and is close to complete indifference. They give a street performer some change partly because of that very weak motive of charity, but mostly because they want to impress their date. Their action seems not to be any good even if they have a motive with an appropriate object because that motive is too weak. Thus for any appropriately oriented motive to be good, it must meet some threshold of strength.[21] Wherever the threshold is, assume that, in all the cases considered in this paper, the motives that I have considered good all

[20] A detailed account is not needed for our purposes, but for a detailed account of the goodness or badness of an attitude as a function of both its strength and the value of its objects, see for instance Hurka (2001).

[21] Susan Wolf has also suggested setting a "minimum level" (Smith 1991, 286 ft.). But Smith rightly notes that a social average will not do. If it is, we can "render a morally mediocre person praiseworthy for her act merely by strategically killing that segment of the population whose morally good desires are stronger than hers" (1991, 286 ft.). Hurka (2001, 81–2) offers an argument without averages for a minimum level. Intuitively, being indifferent or being neither moved to, say, help nor hurt another is positively bad. Now, someone that is even a bit moved by a concern to help is not indifferent. But it seems odd to think that just having infinitesimally more concern than indifference is positively good. Despite a minute change, there is a discontinuity from having a bad motive to having a good one. There should be a gradual shift from indifference to a minimally good motive with a point right before where the motive is purely neutral. In practice, the precise point where a motive is just barely strong enough to be good is difficult to discern. But in principle, it is not when a motive is so weak that it is barely different from indifference.

meet the relevant threshold.

Now, since a motive's value does not depend only on its strength, an action can be good even if its good effective motives are weaker than the bad ones. An action can still be good if its good motive's object is significantly good (right) while its bad one's is trivially bad (wrong), and the latter is not much stronger than the former. Consider two cases.

> *White Lie*: Your friend is having a difficult time right now. They have just been laid off and are showing signs of severe depression. They invite you over to share homemade muffins. Your friend asks if you liked their muffins. You lie and convincingly say that you love them and cannot stop eating them partly because you want to get them out of a depressive spiral, but partly, and a bit more, because you want to tell a little lie.

> *Hard Truth*: Your partner is entering a famous pastry competition that they have dreamed of winning since they were a child, and they have been practicing making different pastries. They give you a soufflé to try. You think it tastes bad. But they will feel bad if you say that. You still say that partly because you want them to improve and accomplish their dreams, but partly, and a bit more, because you want your neighbour, who you dislike, to lose.

First, suppose that, in each scenario, each motive by itself sufficed to prompt its own action. Here, I am inclined to think that, given the scenarios, each of the actions is on balance good. Despite each being prompted partly by a bad motive that is a bit stronger than the good one, the overall goodness of the good motive still outweighs the overall badness of the bad one because the good motive's object seems significantly more good than the bad one's seems bad. In *White*

Lie, wanting to pull your friend out of depression is significantly more good than wanting to tell a lie is bad. As the duty of beneficence far outweighs the duty of honesty here, benevolence is much more good than dishonesty is bad. On the whole, your lying is good even though your bad motive to lie is a bit stronger than your good motive to prevent depression because the goodness of the latter motive outweighs the badness of the former. In *Hard Truth*, wanting your partner to improve their skills and accomplish their dreams is much more good than wanting your neighbour to lose is bad. You should care much more about your partner. So your truth-telling is still good even if its bad motive is a bit stronger than its good one.

The above analysis applies even in cases where neither of your motives sufficed on its own. Suppose you now have a temptation to not do what you actually do, so now need both to act. For example in *White Lie+*, you now also have a strong desire not to have to eat more muffins, and so have a temptation to be honest that they taste bad, and that you want no more of them. So in this new case, your desire to lie and your desire to help your friend are both needed. Similarly in *Hard Truth+*, you now also have a strong desire not to have to try more soufflés, and so have a temptation to lie that they taste great and that your partner has mastered them. So in this new case, your desires for your partner to improve and accomplish their dreams and for your neighbour to lose are both needed. Your effective motives, however, are just as strong in the new cases as they were in the old cases. So your actions in the new cases are still good. And they are so even though their good motives are weaker than their effective bad motives.

3.2 Minimum or Surplus

We see now that an action's value depends on both the strengths and objects of its motives, but how exactly should we weigh the good motives against the other ones? We may think that

we should weigh cases where one has just enough motivation to act distinctly from cases where one has more than enough to act. Ross has suggested this (1930, 170–2). I, however, defend a unified weighing model for both types of cases.

In what I call "surplus cases," your effective motives are more than enough for you to act. Here, Ross thinks that a mixed action's value is the sum of the value of its effective motives. To illustrate, imagine as he does that you may have two motives of the same given strength, one of them is the motive of duty, and the other is either love, pleasure or malice, and that these motives have values of 10, 8, 0, and –8, respectively (1930, 170).[22] Also, imagine that each motive suffices by itself but each actually contributes equally to your action. Ross says that if you act from both motives of duty and love, your action's net value is "10 + 8" or 18 (1930, 171). So in surplus cases, the best action is done from all of the available good motives. If you act from both motives of duty and pleasure, your action's net value is "10 + 0" or 10; (1930, 171). This result captures the idea that if you act partly from a sufficient good motive, acting also from a neutral one does not diminish the action's goodness. He does not discuss acting from good and bad motives where each is self-sufficient, but the same analysis applies: if you act from both motives of duty and malice, your action's net value is "10 − 8" or 2. And this result captures the idea that your action may not be on balance bad but on balance good even if you act partly from a bad motive because you also act from a much better one.

In what I call "minimum cases," your effective motives are just enough to prompt action, and so your effective motives are all needed to prompt it. In such a case, Ross thinks that other things

[22] Ross believes that the desire for your own pleasure is a neutral motive, but others will disagree. Also, he believes that the motive of duty has a greater value than the motive of love, but others will disagree. Despite this potential disagreement, the main points from the analysis still stand.

equal, a mixed action is less good than a pure one done only from a good motive. Imagine again that you are acting partly from duty and partly from love, pleasure or malice. Again, both motives are equally strong and both contribute equally to your action. But now, your motives are "just strong enough to produce [the action]" (1930, 171). Ross suggests that if you act from both motives of duty and love, your action's net value is "$\frac{10+8}{2}$" or 9, which is less than acting solely from that motive of duty (1930, 171). So in minimum cases, he thinks that the best action is done solely from the best motive, which for Ross is duty. And if you act from both motives of duty and pleasure, your action's net value is "$\frac{10+0}{2}$" or 5 (1930, 171). This captures the idea that, with limited motivation, motives compete for effectiveness, and your action is less (more) good when your neutral (good) ones are more effective. And if you act from both motives of duty and malice, your action's net value is "$\frac{10-8}{2}$" or 1 (1930, 171). This captures again the idea that your action can be good even if you need a bad motive to act.

Although Ross suggests that we should split the difference in minimum cases, I argue that we should take a sum like in surplus cases. Note that, in both minimum and surplus cases, Ross uses the same values for the motives even though the strengths of the motives may differ. But a motive's value varies with its strength. If the motives' values are the same in both cases, they must be equally strong in both cases. The only difference is that, in the minimum case, you need both effective motives to act because your countervailing motives must be stronger. But if so, it is not obvious to me that your actions in the minimum case are less good than, and precisely half as good as, yours in the surplus case. And here I am making comparisons only between actions from the same mix of motives, say, from both duty and pleasure. So it seems that, for your actions to differ in value in the two cases, your motives must also differ in value. For your actions in the minimum case to be half as good as their twins in the surplus case, your motives

must be less strong, so your good (bad) motive is less good (bad). (If we imagine that the good (bad) motive's value equals its strength, your motives must be half as strong.[23]) Now, if the difference in the values of actions between the minimum and surplus cases can be accounted for simply by the difference in the values of the motives, we should be able to assess actions in minimum cases by summing up the value of the motives just as in surplus cases.

To see clearer the continuity between the two types of cases, consider two boundary cases. Suppose I would help a student only if my motives to help had at least a total strength of 10. First, in the minimum case, I have a good motive, say, of duty, to help with a strength of 9 and a neutral one, say, of reputation, to help with a strength of 1. I have just enough to help. Second, in the surplus case, my good motive is 1 unit stronger, and my neutral one is the same. I have a bit more than I need to help. Supposing the good motive's value equals its strength, the first action has a value of 9, and the second has a value of 10, ranking higher as expected.[24] So the summative version of the weighing model captures both minimum and surplus cases.

Even if surplus cases are impossible, the model still works; we just never evaluate them. And if surplus is possible but we think only what is needed should count in assessing an action because we adopt a capped view of motivation, where only what you need to act is effective, we can assess the action based on its effective motivation as capped by countervailing motives. If we adopt an uncapped view, we can also assess an action based on its effective motivation. So whether the capped or uncapped view is right, the weighing model still works.

---

[23] There are various ways of conceptualizing the value of motives in terms of their strengths and objects. Linear models are problematic, however, because there should, I think, be diminishing marginal returns. But they suffice for illustration. See Chapter 3 of Hurka (2001) for two better alternatives.

[24] Again, linear models are problematic, but they suffice for illustration.

4 Conclusion

I have defended the weighing model: the goodness of your good effective motives must outweigh the badness of your bad effective motives for your action to be good on the whole. The weighing model captures what the other models have gotten right and improves on them. A good effective motive in proportion to its object and strength makes an action good, and similarly, a bad effective motive in proportion to its object and strength makes an action bad. Neutral motives by themselves, however, make an action neither good nor bad. On the whole, an action is good only if the value of its good motives outweighs the value of its bad motives. Of course, I have defended this claim only as a necessary condition as I am inclined to think that the proportioning of the mixed motives can also make an action bad at least in a respect. Even when the good effective motives outweigh the bad ones, an action can be on balance bad if the neutral effective motives are excessively stronger than the good ones. But it is still true that an action is ever good only if its good effective motives outweigh its bad ones because neutral motives obviously can never make an action any good. So despite this one reservation, we have shed light on weighing mixed actions which should help us make fairer assessments of actions across the various spheres of our lives along a nuanced spectrum of good (and bad). So we may never be saints – still, we act in shades of good.

References

Aboodi, Ron. "One Thought Too Few: Where *De Dicto* Moral Motivation is Necessary." *Ethical Theory and Moral Practice* 20 (2017): 223–237.

Arpaly, Nomy and Timothy Schroeder. *In Praise of Desire*. Oxford: Oxford University Press, 2013.

Arpaly, Nomy. *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford: Oxford University Press, 2002.

Di Norcia, Vincent and Joyce Tigner. "Mixed Motives and Ethical Decisions in Business."*Journal of Business Ethics* 25 (2000): 1–13.

Dubbink, Wim and Luc van Liedekerke. "Rethinking the Purity of Moral Motives in Business: Kant Against Moral Purism." *Journal of Business Ethics* 167 (2019): 379–393.

Henson, Richard. "What Kant Might Have Said: Moral Worth and the Overdetermination of Dutiful Action." *Philosophical Review* 88 (1979): 39–54.

Herman, Barbara. "On the Value of Acting from the Motive of Duty." *Philosophical Review* 90 (1981): 359–82.

———. *The Practice of Moral Judgment.* Cambridge: Harvard University Press, 1993.

Howard, Nathan. "One Desire Too Many." *Philosophy and Phenomenological Research* 102 ([2019] 2021): 302–317.

Hurka, Thomas. *British Ethical Theorists from Sidgwick to Ewing*. Oxford, UK: Oxford University Press, 2014a.

———. "Many Faces of Virtue." *Philosophy and Phenomenological Research* 89 (2014b): 496–503.

———. "Right Act, Virtuous Motive." *Metaphilosophy* 41 (2010): 58–72.

———. *Virtue, Vice, and Value*. Oxford: Oxford University Press, 2001.

Isserow, Jessica. "Moral Worth and Doing the Right Thing by Accident." *Australasian Journal of Philosophy* 97 ([2018] 2019): 251–264.

———. "Moral Worth: Having It Both Ways." *Journal of Philosophy* 117 (2020): 529–26.

Johnson King, Zoë A. "Accidentally Doing the Right Thing." *Philosophy and Phenomenological Research* 100 ([2018] 2020): 186–206.

———. "Praiseworthy Motivations." *Nous* 54 ([2019] 2020): 408–430.

Kant, Immanuel. *Groundwork of the Metaphysics of Morals*. In *Practical Philosophy*, translated and edited by Mary Gregor, 37–108. Cambridge: Cambridge University Press, [1785] 1996.

Markovits, Julia. "Acting for the Right Reasons." *Philosophical Review* 119 (2010): 201–242.

McDowell, John. "Virtue and Reason." *Monist* 62 (1979): 331–350.

Mullin, Amy. "Dependent Children, Gratitude, and Respect." *Journal of Moral Philosophy* 13 (2016): 720–738.

———. "Gratitude and Caring Labor." *Ethics and Social Welfare* 5 (2011): 110–122.

Ross, W. D. *The Right and the Good*. Oxford: Oxford University Press, 1930.

Sliwa, Paulina. "Moral Worth and Moral Knowledge." *Philosophy and Phenomenological Research* 93 (2016): 393–418.

Smith, Holly M. "Varieties of Moral Worth and Moral Credit." *Ethics* 101 (1991): 279–303.

Sorensen, Kelly. “Counterfactual Situations and Moral Worth.” *Journal of Moral Philosophy* 11 (2014): 294–319.

Stratton-Lake, Philip. *Kant, Duty, and Moral Worth*. London: Routledge, 2000.

Stohr, Karen. “Moral Cacophony: When Continence is a Virtue.” *Journal of Ethics* 7 (2003): 339–363.

Sverdlik, Steven. “Kant, Nonaccidentalness and the Availability of Moral Worth.” *Journal of Ethics* 5 (2001): 293–313.

Tomlinson, Laura. “The Moral Worth of Intentional Actions.” *Philosophy and Phenomenological Research* 101 ([2019] 2020): 704–23.