



Better but Wrong: Assessing Conflicts Between the Deontic and the Evaluative

David Faraci¹

Received: 21 October 2024 / Revised: 15 January 2025 / Accepted: 22 January 2025
© The Author(s) 2025

Abstract

In recent work, Benjamin Ferguson and Sebastian Köhler take interest in the claim that permissible acts are always morally better than impermissible acts (BOP). They argue that BOP is both commonsensical and supported by powerful theoretical considerations. They then present a series of cases in which common moral claims appear to conflict with BOP. In this paper, I first show that some of the conflicts Ferguson and Köhler identify are merely apparent, as they arise only given theoretical commitments regarding the value of actions that supporters of the relevant claims can and do reject. In the second half of the paper, I address two other cases Ferguson and Köhler discuss. I argue that the conflicts with BOP in these cases are genuine, but also to be expected, because these cases involve supererogation, and our best theories of supererogation are in tension with Ferguson and Köhler's case for BOP.

Keywords Permissibility · Reasons · Value · Deontology · Supererogation

Benjamin Ferguson and Sebastian Köhler take interest in the claim that “permissible acts are morally better than impermissible acts.”¹ More formally:

Betterness of Permissibility (BOP) For all actions x and y , where x and y are elements of a set of options available to an agent in a particular context, if x is permissible and y is impermissible, then x is morally better than y .²

¹ Unless otherwise noted, all citations are to Ferguson and Köhler (2019).

² This is a paraphrase of: “Let X be a finite set of (token) acts, Z be the set of all nonempty subsets of X , and A be an agenda [“the set of options available to an agent in a particular context”], such that $A \in Z$. . . [F]or all agendas A in Z and all acts x, y in A , if x is permissible and y is impermissible, then x is morally better than y ” (2454).

✉ David Faraci
david.n.faraci@durham.ac.uk

¹ Department of Philosophy, Durham University, 51 Old Elvet, DH1 3HN Durham, United Kingdom

Ferguson and Köhler argue that BOP is both commonsensical and supported by powerful theoretical considerations. They then present a series of cases in which common moral claims appear to conflict with BOP.

In this paper, I first show that some of the conflicts Ferguson and Köhler identify are merely apparent, as they arise only given theoretical commitments regarding the value of actions that supporters of the relevant claims can and do reject. The lesson of Ferguson and Köhler's discussion of these cases, I argue, is that as we move towards more ecumenical application of goodness—one which goes beyond states of affairs—we need to take care to avoid importing consequentialist assumptions about the potential grounds for goodness.

In the second half of the paper, I address two other cases Ferguson and Köhler discuss. I argue that the conflicts with BOP in these cases are genuine, but also to be expected, because these cases involve supererogation, and our best theories of supererogation are in tension with Ferguson and Köhler's case for BOP.

1 Consequentialism about Action Value

Ferguson and Köhler recognize that consequentialists may object to the idea that acts, rather than merely their outcomes, can be good. Especially since they wish to remain neutral between moral theories, they feel pressed to address this worry, which they do as follows:

We will be concerned primarily with moral goodness as a property of *acts* and do not follow the more common restriction of goodness to states of affairs. Admittedly, this is already a significant assumption, since it conflicts with the standard consequentialist claim that moral goodness is a property of states of affairs. By focusing on the goodness of acts, we do not deny that states of affairs can also be good. And indeed, we think consequentialists can easily endorse BOP. Consequentialism is compatible with BOP if acts' goodness derives from, or is grounded by, the fundamental (for consequentialists) goodness of state of affairs. (2453)

Call this consequentialist-friendly idea, that the goodness of acts derives from the goodness of states of affairs, Consequentialism about Action Value (CAV). In this section, I argue that two of the common moral claims Ferguson and Köhler discuss *only* conflict with BOP if some version of CAV is true. Since rejecting (that version of) CAV is well-motivated for supporters of the relevant claims, I conclude that these cases are not well-positioned to illustrate tension between common moral claims and BOP.

First, consider Ferguson and Köhler's discussion of the doctrine of double effect (DDE):

BOP is incompatible with an interpretation of the doctrine of double effect, which claims "there is a moral constraint on intending evil (such as harm), even when the evil will be a means to a greater good" (Kamm, 2007, p. 21). . . [A]lthough BOP is not directly incompatible with deontology, it is incompat-

ible with permissibility assessments that differ from teleological permissibility assessments, such as those that follow from the doctrine of double effect. Deontologists probably have the strongest theoretical reasons to resist BOP. (2459)

Ferguson and Köhler's position seems to be that the DDE is inconsistent with BOP because it entails that intending evil is always impermissible, yet sometimes morally better than not intending evil, which is permissible. But Kamm doesn't say that intending evil is ever better than not doing so; she says only that it is sometimes a *means* to a greater good. Accordingly, the DDE only conflicts with BOP if we accept CAV—specifically, a version of CAV linking actions' value to the value of the states of affairs to which they are a means. Yet the clear deontological bent of the DDE, on the one hand, and consequentialist bent of CAV, on the other, mean we should expect theorists who endorse the former to reject the latter.

In fact, some nonconsequentialists have already argued against CAV. Consider, for example, Paul Hurley's (2017) response to the common idea that consequentialism enjoys strong intuitive support from the so-called Compelling Idea that it is always morally permissible to do whatever will lead to the morally best outcome.³ Hurley argues that this seems compelling because it is implied by two further claims. The first, which he agrees is compelling, is that acts can be good or bad, and that it is always permissible to do whatever act is morally best. He calls this the General Idea. The second idea, which he finds unconvincing, is CAV.

Here is another of Ferguson and Köhler's cases:

Suppose both A and B are equally good candidates for a kidney transplant, but only one kidney is available. The following three claims jointly violate BOP:

- (L1) the moral value of A receiving the kidney is equal to the moral value of B receiving the kidney (because, say, the welfare gains under each outcome are equal).
- (L2) there are moral constraints that apply to how the kidneys are allocated: giving a kidney *directly* to either A or B is impermissible. It is only permissible to use a fair coin toss to select the recipient.
- (L3) the moral value of a randomised act is equal to the expected moral value of its parts (and you think the 'parts' here are simply the act of giving a kidney to A and giving a kidney to B). (2458)

In conjunction with L2, L3 entails that choosing the recipient on non-random grounds is just as good as choosing on random grounds. This generates inconsistency: BOP entails that the members of an equally good pair of actions cannot be permissible and impermissible, respectively; L1 and L2 entail that some are.

The culprit is L3. For nonconsequentialists, the point is much the same as with the DDE: nonconsequentialists are likely to take things other than an action's

³ See, e.g., Scheffler, 1982.

consequences—such as whether it was taken intentionally or as “part” of a randomized action—to be relevant to its evaluation, and thus to reject L3.

But even consequentialists needn’t accept L3. Certainly, actualists who take an action’s value to be a function of its actual outcome needn’t. But expected value consequentialists needn’t, either. Here is one natural way of avoiding this result: It has become standard for consequentialists wishing to accommodate evaluative grounds classically associated with nonconsequentialism to accept that an action’s consequences include the fact that that action was taken. This allows them to maintain, for instance, that killing is worse than letting die, because only the former has the consequence *that killing has taken place*.

A consequentialist might similarly reject L3 by maintaining that *someone’s having chosen who lives and dies* is, all else equal, a bad-making feature of a state of affairs. This would allow them to avoid conflict with BOP without rejecting CAV outright: They needn’t deny that the value of actions depends on the value of states of affairs; they merely need to deny that the relevant states of affairs are limited to those that would support L3.

We’ve seen that supporters of the DDE and randomized kidney allocation can avoid conflict with BOP by denying CAV—something some actual nonconsequentialists have already done. We’ve also seen that consequentialists who support randomized kidney allocation can avoid conflict with BOP by using a standard strategy for accommodating evaluative grounds classically associated with nonconsequentialism, and thus rejecting a particular version of CAV.

I conclude that these cases are not well-suited to illuminate tension between common moral claims and BOP. What Ferguson and Köhler need are cases where people make direct judgements about the value of an action that don’t align with their judgements about reasons or permissibility, rather than about the value of a controversial potential grounds for the value of an action. We’ll look at two such cases in § 3. First, I briefly discuss Ferguson and Köhler’s theoretical case for BOP. This serves to reinforce the connection between their discussion and worries like Hurley’s, as well as set things up for § 3.

2 The Case for BOP

Ferguson and Köhler’s main argument for BOP is that it is entailed by two further claims:

Weak Classical View For all actions x and y , “there is at least as strong a moral reason to x as there is to y , if and only if x is at least as morally good as y .”

Reasons Monotonicity of Permissibility For all actions x and y , “if x is permissible and the moral reasons for performing y are at least as strong as those for performing x , then y is also permissible.”

The Weak Classical View ties moral reasons for actions to the moral goodness of those actions, while Reasons Monotonicity says that permissible actions are always better supported by moral reasons than impermissible ones. Together, these imply

that permissible actions are always morally better than impermissible ones—i.e., BOP.

There is no direct connection between either of these claims and CAV. However, since the Weak Classical View concerns the goodness of actions, one might wonder whether Ferguson and Köhler's discussion of that claim encourages us to accept CAV.

Most of their discussion does not bear on CAV at all. This is as expected, given their explicit aim to remain neutral between moral theories. In keeping with that aim, when they discuss potential problems for the Weak Classical View, they don't say anything about the *nature* of the goodness of actions; they simply address the possibility that the goodness of actions—whatever that is—could be unnecessary or insufficient for moral reasons.

The one relevant part of Ferguson and Köhler's discussion addresses the possibility that supporters of apparently BOP-inconsistent claims might reject the Weak Classical View by claiming that rightness grounds goodness:

- If rightness is the common source of both reasons and moral goodness, then all acts we have moral reason to perform (because they stem from rightness) will still be morally good (because they also stem from rightness). . . . Thus, a rightness-based objection to the weak classical view must show not only that (1) rightness can be the source of moral reasons, but also (2) some acts that rightness-based considerations provide us with moral reasons to perform are not morally good in any respect. (2461)

However, Ferguson and Köhler object to such rightness-first views:

- Since permissibility is a modal concept that has only two possible values [permissible or impermissible], if moral reasons are a function of permissibility, then moral reasons can also only take one of two values. We would either have a moral reason to perform an action or we would not. . . . However, the weak classical view understands reasons differently, as considerations that come in varying degrees. It is a claim not about what we have all things considered moral reasons to do, but about the association of 'smaller', *pro tanto* moral reasons with moral goodness. (2461)

To an extent, this is all beside the point: Supporters of the DDE or randomized kidney allocation don't need to reject the Weak Classical View. They just need to deny that bringing about a greater good makes an act more valuable, and that having equal expected value makes acts equally valuable. That said, it does seem likely that nonconsequentialists will need *some* story about the grounds for the value of actions—some alternative to CAV. If there were no viable alternatives, that would certainly be a problem!

But, of course, Ferguson and Köhler haven't shown that there are no viable alternatives; they have merely offered an objection to one alternative. And we have no reason to doubt there are others: the goodness of an action can be grounded in *any* feature of that action, not just its consequences.

It is not hard to see, however, why the alternative they discuss might seem like the default: it is suggested by the classical dichotomy between consequentialist views according to which the Good precedes the Right, and deontological views on which the Right precedes the Good. Given this, I believe what Ferguson and Köhler's points here most importantly do is support another claim of Hurley's: that wide acceptance of this dichotomy has been an unfortunate thumb on the dialectical scales in consequentialism's favor:

Ironically, opponents of consequentialism have been complicit in eliding from view the alternative, less consequentialism-friendly ... specification of the General Idea. Consequentialism has often been presented by both its critics and defenders as holding that the Good is prior to the Right. Deontological alternatives are characterized, by contrast, as holding that the Right is prior to the Good. In one fell swoop, such a formulation of the debate effectively relegates the central role of telic evaluation in normative ethics to outcomes, and commits the deontologist to holding that the right action is determined prior to, and independently of, the relevant consideration of goodness. The effect of such a framing is to elide from view evaluation of actions and reasons for action as good, better, and best, ceding any specification of the "best" in the intuitive General Idea to the evaluation of outcomes. (Hurley, 2017, p. 34)

Hurley warns that by assuming all goodness is a feature of outcomes, we make consequentialism look like the only game in town. I think Ferguson and Köhler are best read as driving Hurley's point home, showing just how confused common moral claims like the DDE can look when widespread, implicit reliance on CAV goes unrecognized, and when we unjustifiably assume that we must either ground rightness in the goodness of outcomes, or ground goodness in the rightness of actions. Ultimately, the lesson we should draw from their discussion of these cases is not that some common moral claims are inconsistent with BOP, but rather that as we move away from talking about goodness as a feature of outcomes only, we need to take care not to implicitly import consequentialist-leaning assumptions.

3 BOP and Supererogation

Turn now to a third case: Barbara is considering opening a factory in an impoverished country, where she would be able to her to pay employees an unfair wage that is nevertheless significantly higher than any alternative they have. Ferguson and Köhler point out that many would accept that it is:

- (E1) Permissible to not engage in mutually beneficial transactions, but
- (E2) Morally better to engage in mutually beneficial transactions than to not,

(E3) Impermissible to gain unfairly in mutually beneficial transactions. (2457)

E1 and E3 suggest that while Barbara's not opening a factory is permissible, paying her employees unfairly is impermissible. E2 suggests that the latter is better than the former. This is inconsistent with BOP.⁴

Some will be inclined to reject E1 or E3. Perhaps Barbara is obligated to enter into this arrangement, since it maximizes utility (e.g., Bailey, 2010). Or perhaps it is permissible for her to exploit her workers given that this involves consensual benefits (e.g., Powell & Zwolinski, 2012). But suppose you are inclined to accept E1 and E3. You could then instead pursue the same strategy used in the first two cases: You could accept that the outcome of exploitation is better than the outcome of doing nothing, but deny that this means that the act of exploitation itself is better than non-transaction, and thus deny E2.⁵

There is nothing wrong with this move *per se*, but it faces an important challenge. Suppose Barbara did not have the option of paying her employees a living wage, that she was limited to either paying them a low wage or not hiring them at all. In that case, surely the former would be the better thing to do, lending intuitive force to E2 itself—a claim about the value of an action—and not merely to a claim about the value of its outcome. There is no equivalent move in the other cases. A forced choice between intending harm or forgoing a greater good would not move anyone attracted to the DDE to claim that the former is better. A forced choice between selecting the kidney recipient at random or using some other method would not move anyone attracted to L2 to claim that the two are equally good. Indeed, stipulating a forced choice adds nothing to these cases at all.

Barbara's case is tricky. For instance, one might reply that in the forced choice case, paying a low wage is indeed the better thing to do, but that this is irrelevant because what's impermissible in the original case isn't paying a low wage, but rather paying an exploitative one, where exploitation depends on the existence of the option to pay more.

Luckily, Ferguson and Köhler open their article with a case that more cleanly illustrates the relevant sort of tension with BOP, a relative of Horton (2017)'s so-called All or Nothing Problem for supererogation:

Suppose that as you stroll along a beach one evening a storm suddenly strikes. You soon hear the panicked shouts of two swimmers who will surely not last long in the growing swells. Though you'd like to help them, you're no seaman and the conditions are extremely perilous; if you put out to save them, it is possible you'll all drown. We think most people would agree that, given the risk, you are not morally obliged to save the swimmers. However, it also seems that if you could save even one, this would be morally better than staying on the beach. Nevertheless, suppose that you spot a dinghy on the beach and heroi-

⁴ In earlier work, Ferguson (2016) refers to this as the "paradox of exploitation," echoing the more general "paradox of deontology."

⁵ I discuss related mismatches between deontic and evaluative judgements in cases of wage exploitation in Faraci (2019).

cally attempt a rescue. When you reach them, the swimmers are side-by-side and you could pull both into the boat and make it back to shore. However, suppose you save only one, leaving the other to drown. We think most people would agree that if you do put out to save the swimmers, given that you could save both, it is impermissible to save only one. (2452)

The story suggests that while it is permissible for you not to attempt a rescue and impermissible to save only one swimmer, the latter is better than the former.

As described, this case is actually not inconsistent with BOP, because BOP is a claim about the relationship between acts within option sets, and there are two option sets here. The first option set is: (a) stay on shore or (b) attempt a rescue. Suppose you attempt a rescue, make it out to sea and get one of the swimmers into the boat. Then the second option set is (c) pull the other swimmer into the boat or (d) simply head back to shore, leaving the second swimmer to drown.⁶

But this is a red herring, because we can alter the case such that many people would continue to make BOP-inconsistent judgements relative to a single choice set. Suppose that, instead of a dinghy, there is a robotic boat on the shore, with the following features:

- (a) if you get into the boat and press the button labeled 'Both', the boat will take you out to the swimmers and automatically rescue both of them, assuming it does not capsize;
- (b) if you get into the boat and press the button labeled 'One', the boat will take you out to the swimmers and automatically rescue exactly one of them, assuming it does not capsize (and you will be unable to do anything to help the other);
- (c) if you do not get into the boat and press a button, both swimmers will drown;
- (d) there is a non-negligible chance that the boat will capsize—the same chance regardless of which button you press—in which case everyone will drown.⁷

⁶ We can see why it is important to limit BOP to a claim about option sets by noticing that even classical, maximizing utilitarianism would otherwise be inconsistent with BOP. Suppose I do something that produces 30 utils, which is obligatory (and therefore permissible) because it maximizes utility. Later, I am presented with two options which would produce 40 and 60 utils, respectively. Obviously, it would be absurd to conclude that it must be permissible for me to produce 40 utils, because this is morally better than producing 30 utils, and producing 30 utils was permissible. Applying this to the drowning case: Suppose we treat each death as -1 value. The expected value of staying on shore would be greater than that of attempting a rescue if the probability that you will fail to rescue the swimmers and drown yourself is greater than 2/3. (At 2/3 the expected values are equal: -2 for staying on shore, since the swimmers will drown; $1/3 \cdot 0 + 2/3 \cdot -3 = -2$ for a rescue attempt.) If you ignore the risk and manage to successfully rescue one swimmer, the fact that "the swimmers are side-by-side" implies that the second rescue involves no additional risk. This means that with respect to your new option set, the expected values of attempting to rescue the second swimmer versus leaving the second swimmer to drown are 0 and -1, respectively, making the latter impermissible. This is all consistent with BOP. But this consistency is a red herring, as I discuss momentarily, because we can adjust the case to involve a single option set, and most people's judgements will then violate BOP because they will regard the permissible action of staying on shore (-2) as better than the impermissible action of leaving the second swimmer to drown (-1).

⁷ Thanks to Tristram McPherson for this case, and for using it to illuminate for me the redness of this herring.

As in the other cases, it is possible to render these judgements consistent with BOP by maintaining that while the survival of one is better than the survival of neither as a state of affairs, pressing One is not a better action than remaining on shore. However, as in Barbara's case, there is independent intuitive motivation for the claim about action value: if there were no Both button, surely pressing One would be the better thing to do.

It is no coincidence that this case is a relative of a puzzle about supererogation. This is because, as I now argue, the cases Ferguson and Köhler identify where there is conflict with BOP even once we reject CAV—including the drowning case and some readings of Barbara's—rely on a tension between Ferguson and Köhler's case for BOP and our best theories of supererogation.

To begin, recall the second claim Ferguson and Köhler use to motivate BOP:

Reasons Monotonicity of Permissibility For all actions x and y , “if x is permissible and the moral reasons for performing y are at least as strong as those for performing x , then y is also permissible.”

The monotonicity condition paints a particular picture of the deontic: imagine a line along which all of an agent's options are ordered, from left to right, by the degree to which they are supported by moral reasons; somewhere along the line is a *permission* point; everything at or to the right of that point is permissible; everything to the left of that point is impermissible. If staying on shore is permissible but pressing One is impermissible, the latter must fall to the left of the former. But this is implausible; surely there are weightier moral reasons to save the one than to save neither; that's precisely why that seems like the better thing to do.

Current theories of supererogation provide a diagnosis. Notice that in the drowning case, intuitively, it's not that you are permitted to remain on shore because that is sufficiently supported by the combined weight of all moral reasons; rather, you are permitted to remain on shore because you aren't morally required to risk your life, even to do what the weight of reasons most supports. Precisely because of intuitions like this, the best contemporary theories of supererogation all reject the one-dimensional model of reasons that entails the monotonicity condition, a model on which the deontic status of an act is a simple function of the weight of the moral reasons that support it. For some theorists, the risk to your life is still just a moral reason, but talk of the “strength” or “weight” of moral reasons is not univocal; reasons contribute to deontic status in different ways (e.g., Archer, 2016; Lazar, 2019; Parfit, 2011). For others, the risk to your life is not just a moral reason; it grants a moral *prerogative* that justifies independently of moral reasons (e.g., Hurka & Shubert, 2012; Kamm, 1985; Muñoz, 2020; Scheffler, 1982; Slote, 1984). Either way, the picture painted by the monotonicity condition is too simplistic. As Daniel Muñoz pithily puts it:

These are the leading solutions in the literature, and I can't stress enough that they both involve the idea that right and wrong depend on two dimensions. There isn't just a cutoff along a single scale from naughty to nice. (Muñoz, 2020)

- Structurally, the two-dimensional model is consistent with the monotonicity condition. But many of the intuitions about cases involving supererogation that motivate the model are not, as Ferguson and Köhler have shown. That might seem to take us back to square one: do we give up the intuitions about cases or the monotonicity condition? But as I'll argue, the case for the monotonicity condition implicitly relies on the one-dimensional model. We thus face a choice between common intuitions and an apparently viable two-dimensional model of reasons that captures them, on the one hand, and a condition motivated by the one-dimensional model of reasons, which is inconsistent with our intuitions and for which no argument has been offered, on the other. That looks like an easy choice in favor of our intuitions already, and we can cap it with Joshua Gert's debunking explanation for wide implicit acceptance of the one-dimensional model:
- Philosophers have certainly asked themselves about the role that normative reasons play in determining [deontic] status. The problem is that they generally have given only the most abstract and formal of answers. . . . [T]hey do nothing to indicate what "counting in favor of" or "supporting" amount to, and they certainly make it extremely unlikely that different ways of counting in favor or supporting will be distinguished. . . . Because of this most philosophers who make use of the notion of "the strength of a reason" have been able to use that notion only in a metaphorical way. The natural metaphor here is with the strength of a physical force such as weight. . . . As a result, these philosophers have thought that in contests between two opposed reasons, they could discern the stronger reason as the one that rationally compels action, and thereby "wins," and have supposed that they could define [right] action as "the one supported by a preponderance of reasons." (Gert, 2007, 535–36).
- To see how the case for the monotonicity condition relies on the one-dimensional model, let's return to Ferguson and Köhler:

If the condition is false, then there are some acts that are impermissible, but for which there are stronger *moral* reasons to perform than some permissible acts. Now, *prima facie*, this implication seems to conflict with the roles these deontic concepts play in answering questions about what to do. (2465)

Note the framing here in terms of a univocal "strength" of reasons, as per Gert's comments. Predictably, our intuitions about the drowning case only "conflict with the roles these deontic concepts play in answering questions about what to do" if the strength of reasons is univocal in this way. To see this, consider a toy two-dimensional model on which reasons contribute to deontic status in two ways: they have requiring weight and they grant prerogatives. Reasons' requiring weight combines to determine a ranking of options from "naughty to nice," as on the one-dimensional model. Deontic status is a function of this ranking (and a "permission point") except where reasons grant prerogatives, intervening to

directly render certain actions permissible—actions which can fall anywhere on the weighted scale.⁸

On this model, the reasons for pressing One have greater requiring weight than those for remaining on shore, but the risk to your life grants a prerogative to remain on shore. At the same time, pressing One is impermissible because the reasons for pressing One have *less* requiring weight than (indeed, they are a proper subset of) those for pressing Both. And there are no considerations that could plausibly grant a prerogative to press One rather than Both, especially given that, as stipulated, the risk associated with pressing One or Both is the same. (In the original case, the claim that “the swimmers are side-by-side” similarly implies that saving the second swimmer exposes you to no additional risk.)

Similar points apply to Barbara’s case. On the view Ferguson and Köhler target, hiring workers at an unfair but still beneficial wage is impermissible, while not opening a factory at all is permissible. Yet the claim that it is “[m]orally better to engage in mutually beneficial transactions than to not” seems to imply that the reasons for the former have greater requiring weight than the reasons for the latter. Again, the two-dimensional model eliminates any apparent conflict with the deontic role of reasons: it is Barbara’s prerogative not to open a factory at all, but if she does open a factory there is no prerogative to exploit her workers rather than pay them a fair wage.⁹

Ferguson and Köhler continue:

To find that action *y* is impermissible, but *x* is permissible, is clearly to find that one ought to do the second, but not the first act (provided these are your only two options). It is hard to make sense of this however, if the impermissible act turns out to be the one supported more strongly by the moral reasons. (2465)

Crucially, we haven’t actually seen a case in which an agent has only two options, yet the impermissible one is “supported more strongly by the moral reasons.” Nor will we: on our two-dimensional model, it is always permissible to do whatever is most supported by the requiring weight of reasons, even though one may enjoy a

⁸ Those familiar with the relevant literature will note that this model marries Gert’s “requiring reasons” with Muñoz and others’ “prerogatives.” I do this for two reasons. The first is simplicity: a direct adaptation of Gert’s view, on which requiring weight would be contrasted with justifying weight, would open the door to more complex interactions between the dimensions. This is as it should be; after all, even if you are not required to risk your life to save two, you might be to save a million. But thankfully these complications don’t arise in the cases under consideration, so treating prerogatives as absolute here is a harmless simplification. (Thanks to Daniel Muñoz for discussion on this point.) The second reason is dialectical: I find talk of prerogatives helpfully intuitive, but I also want the model to be “reasons-only” because, as an anonymous referee points out, Ferguson and Köhler stipulatively reserve ‘reason’ for all considerations that contribute to deontic status.

⁹ In the drowning case, it is easy to see how the risk to your life might play a different role from the lives of the swimmers, the former justifying inaction even while the latter determine which action to take, should you take any. But cases like Barbara’s are sometimes framed in ways that obscure this, suggesting that the same considerations that make it permissible for Barbara not to open a factory at all likewise make it permissible for her not to pay a fair wage. I argue that this relies on a mistaken understanding of the considerations at play in Faraci (2019).

prerogative to do something else. This is reflected in points already made about our cases. If pressing Both were not an option, then you would have a choice between pressing One and remaining on shore, and the former would seem the better thing to do. Of course, you would still enjoy a prerogative to remain on shore, but pressing One would now be supererogatory, not impermissible. Similarly, suppose paying a fair wage is not an option and Barbara faces a choice between paying a low wage and not opening a factory at all. Here, again, the former might seem the better thing to do. Again, Barbara would still enjoy a prerogative not to open a factory at all, but paying a low wage would be permissible.¹⁰

Ferguson and Köhler conclude:

The monotonicity condition can be false only if there are considerations other than moral reasons that determine what we ought morally to do. But this is just what reasons are—those considerations that serve the function of explaining and guiding what we ought to do are called reasons. (2465)

The preceding explains why the first sentence is true only given a one-dimensional model. For the monotonicity condition can *also* be false if moral reasons determine what we do, but do so in different ways—e.g., through both having requiring weight and granting prerogatives.

4 Conclusion

Ferguson and Köhler argue that a number of common moral claims are inconsistent with BOP, which is entailed by two theoretical principles: the Weak Classical View and Reasons Monotonicity of Permissibility. I have argued that in some of these cases, the conflict depends on controversial further assumptions about the relationship between the value of actions and the value of states of affairs—assumptions supporters of the claims in question reject. The true lesson of these cases, I claim, is that as we increase our attention to the goodness and badness of things beyond states of affairs—including especially of actions—we must take care not to implicitly retain consequentialist-leaning assumptions.

In the remaining cases, there is genuine conflict with BOP, which can be traced to the fact that common, supererogation-involving intuitions about these cases seem best captured by two-dimensional models that are inconsistent with the one-dimensional model that motivates the monotonicity condition.

Given my arguments in the previous section, it may be tempting to conclude that we should reject BOP. My own view, however, is that this should remain an open question. Gert argues that consistent metaphorical talk of the “weight” of reasons

¹⁰ It feels a bit strained to say that this would be supererogatory. The explanation is not hard to find: we typically reserve that label for selfless acts, but we do not imagine Barbara to, say, be giving up her dream of becoming a farmer to open a factory so as to give what little help she can to the world's poor; we imagine her to have decided to open a factory in an impoverished country, at least in part for self-interested reasons, and then to have found that she can only afford to pay a worryingly low wage.

has discouraged development of multi-dimensional understandings of how reasons issue in deontic verdicts. But another way of putting my points in the previous section is that the only cause we've located thus far for moving away from this kind of simplicity is the complications created by supererogation. Supererogation is widely accepted, but hardly uncontroversial. My own view—which for obvious reasons I won't fully develop here—is that there are at least as good debunking explanations for common belief in supererogation as for the common assumption of one-dimensionalism about reasons.¹¹

Ultimately, therefore, I conclude that whether we should reject two-dimensionalism or BOP depends on two further issues: whether supererogation is the only reason we need two-dimensionalism; and whether we should believe in supererogation. This fits with a lesson I find myself learning over and over in ethics, that little progress is made purely on the basis of formal considerations, because there will almost always be a way to make an intuitive position internally consistent. Here, as in most place, we are forced back to debating the intuitions themselves.

Acknowledgements Thanks to Benjamin Ferguson, Tristram McPherson and Daniel Muñoz for helpful comments.

Declarations

Competing Interests The author reports no competing interests regarding this work and its funding.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Archer, A. (2016). Moral obligation, self-interest, and the transitivity problem. *Utilitas*, 28(4), 441–464.
- Bailey, A. D. (2010). The nonworseness claim and the moral permissibility of better-than-permissible acts. *Philosophia*, 39(2), 237–250.
- Broome, J. (1990). 'Fairness'. *Proceedings of the Aristotelian Society* 91: 87–101.
- Diamond, P. (1967). Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: Comment. *Journal of Political Economy*, 75(5), 765–766.
- Faraci, D. (2019). 'Wage exploitation and the nonworseness claim: Allowing the wrong, to do more good.' *Business Ethics Quarterly*, 29(2):169–188.
- Ferguson, B. (2016). The paradox of exploitation. *Erkenntnis*, 81(5), 951–972.

¹¹ For one, structurally there is little difference between the idea that morality permits us to do less than what's best and the competing possibility that our nonmoral reasons simply outweigh our moral ones in some cases. Yet it is obviously attractive in potentially biasing ways to say that you're still doing what's right, rather than reasonably doing what's wrong.

- Ferguson, B., & Sebastian, K. (2019). 'Betterness of permissibility'. *Philosophical Studies*.
- Gert, J. (2007). Normative strength and the balance of reasons. *Philosophical Review*, 116(4), 533–562.
- Horton, J. (2017). The all or nothing problem. *The Journal of Philosophy*, 64(2), 94–104.
- Hurka, T., and Esther Shubert (2012) 'Permissions to do less than best: A moving band'. In *Oxford Studies in Normative Ethics*, edited by Mark Timmons, 2:1–27. Oxford University Press.
- Hurley, P. (2017). Why consequentialism's compelling idea is not. *Social Theory and Practice*, 43(1), 29–54.
- Kamm, F. (1985). *Morality, mortality, volume II: Rights, duties, and status*. Oxford University Press.
- Kamm, F. (2007). *Intricate ethics: Rights, responsibilities, and permissible harm*. Oxford University Press.
- Lazar, S. (2019). Accommodating options. *Pacific Philosophical Quarterly*, 100(1), 233–255.
- Muñoz, D. (2020). 'Three paradoxes of supererogation'. *Noûs*, 55, 699–716.
- Parfit, D. (2011). *On what matters*. Oxford University Press.
- Powell, B., & Zwolinski, M. (2012). The ethical and economic case against sweatshop labor: A critical assessment. *Journal of Business Ethics*, 107(4), 449–472.
- Scheffler, S. (1982). *The rejection of consequentialism*. Oxford University Press.
- Slote, M. (1984). Morality and self-other asymmetry. *The Journal of Philosophy*, 81(4), 179–192.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.