

## COMMENTS AND CRITICISM

## ALL OR NOTHING, BUT IF NOT ALL, NEXT BEST OR NOTHING\*

Consider the following case from a wonderful paper by Joe Horton.<sup>1</sup>

*Armsgiving*: Two children face a deadly threat. You can either (i) do nothing, (ii) save one child by sacrificing your arms, or (iii) save both children by sacrificing your arms.<sup>2</sup>

It seems plausible that:

- (1) It is permissible for you to do nothing.
- (2) It is wrong for you to save only one child.

Horton observes that, “(1) is plausible because of the sacrifice that saving the children requires” and “(2) is plausible because saving both children requires no greater sacrifice than saving only one,” but “there may be a problem with accepting both of these claims.” “Suppose that you are a bad person, and you dislike one of the children. You are willing to save the other child, but you are not going to save both.” Horton says that (1) and (2) seem to imply:

- (3) “You ought to save neither child rather than save only one.”<sup>3</sup>

Horton then says that (3) seems implausible, briefly elaborating: “Surely the best moral view would not discourage you from saving the one child.”<sup>4</sup>

\*I owe thanks to many for helpful discussions, including audience members at the University of Cambridge, the University of Edinburgh, the London School of Economics, the University of Oxford, the University of St Andrews, Stockholm University, the University of Warwick, and the University of York. For written comments, I am very grateful to Christian Barry, Krister Bykvist, Daniel Muñoz, Tom Sinclair, Travis Timmerman, and the editors of and referees for this JOURNAL. My greatest debts are to Joe Horton and Jeff McMahan, for inspiration, discussion, and comments.

<sup>1</sup>Joe Horton, “The All or Nothing Problem,” this JOURNAL, CXIV, 2 (February 2017): 94–104.

<sup>2</sup>Cases of this sort date back to Derek Parfit, “Future Generations: Further Problems,” *Philosophy and Public Affairs*, xi, 2 (Spring 1982): 113–72, at p. 131.

<sup>3</sup>Following Horton, I will write as though what you *ought* to do is what you are *obligated* to do. We might instead claim that you are not always obligated to do what you ought to do (or have most reason to do). See Justin Snedegar, “Reasons, Oughts, and Requirements,” in Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 11 (Oxford: Oxford University Press, 2016), pp. 155–81.

<sup>4</sup>All of the quotations above are from Horton, “The All or Nothing Problem,” *op. cit.*, p. 94.

In sum: (1) seems plausible, (2) seems plausible, and yet (1) and (2) seem to imply (3), which in turn seems implausible. This is one instance of the *All or Nothing Problem*.<sup>5</sup> I will focus on this instance of the problem, but my discussion below generalizes.

In section I, I very briefly outline Horton's solution to the All or Nothing Problem, noting two difficulties it faces. In section II, I argue that in order for (3) to seem implausible in the way Horton claims it does, (3) must be interpreted so as to imply a conditional-ought claim that would discourage you from saving a child if you are not going to save both children but are willing to save one. In section III, I argue we should reject the principle Horton proposes to bridge (1) and (2) to (3). Instead, we should accept what I call *conditional permissions*. In section IV, I summarize my solution to the All or Nothing Problem, which, I believe, shows what the problem was all along.

#### I. HORTON'S SOLUTION

Horton's solution is based on the claim that, given your willingness to save one child by sacrificing your arms, you cannot reasonably appeal to having to sacrifice your arms to justify saving neither, or only one. He writes:

If you were not willing to save either child, it would be permissible for you not to save either, but because you are willing to save one, you ought to save both.<sup>6</sup>

He goes on to point out that if "we reject (1) in favor of [the above claim], accepting (2) does not commit us to (3). We can instead accept the following claim":

Because you are willing to save one child, you ought to save both, but if you are not going to save both, you ought to do the next best thing, which is to save one. That is, you ought to save one child rather than save neither.<sup>7</sup>

This, in a nutshell, is Horton's solution. It faces the following two difficulties.

<sup>5</sup> See *ibid.*, p. 94, for Horton's general statement of the problem. I take the All or Nothing Problem to be distinct from the problem of what to say about the case in which you are not going to save both children, and refrain from saving either—or make yourself unwilling to save either—merely in order to avoid saving only one (rather than in order to keep your arms). For relevant discussion, see my "Whether and Where to Give," *Philosophy and Public Affairs*, XLIV, 1 (Winter 2016): 77–95, at pp. 89–91; and Horton, "The All or Nothing Problem," *op. cit.*, pp. 97–98. For discussion of holding motives constant in cases like Armsgiving, see Douglas Portmore, *Opting for the Best* (New York: Oxford University Press, forthcoming), section 6.4.

<sup>6</sup> Horton, "The All or Nothing Problem," *op. cit.*, p. 97.

<sup>7</sup> *Ibid.*

First, as Horton is aware, his solution implies that no one ever performs supererogatory acts of beneficence. Take any act that involves a sacrifice to benefit others. In order for one to have performed this act, one had to have been willing to make the sacrifice. But, according to Horton's solution, if one were willing to make the sacrifice, then it would not have been permissible for one not to. Therefore, no performed act that involves a sacrifice to benefit others is supererogatory. This may seem an implausible conclusion. Still, it may be an acceptable conclusion, as it leaves room for *an* account of supererogation (which operates at the level of what sacrifices one is willing to make).<sup>8</sup>

Second, it is controversial that willingness to make a sacrifice is relevant to the wrongness of not making this sacrifice in the way implied by Horton's solution.<sup>9</sup> In cases in which you are unwilling to save *either* child, his solution implies (1), which, together with (2), seems to imply (3). But Horton claims that, if you are unwilling to save either child, (3) is acceptable, because then it does "not discourage anyone who is willing to save one child from doing so."<sup>10</sup> On the other hand, in cases in which you are willing to save at least one child (but not necessarily both), Horton's solution avoids (3) by rejecting (1). Even if (3) is acceptable in cases in which you are unwilling to save either child, rejecting (1) in cases in which you are willing to save at least one may not seem quite right. In some of the latter cases, it may seem particularly implausible to reject (1). Perhaps, since the children are in such obvious peril, you are willing to sacrifice your arms to save them, and yet, since you love your arms so much, you are simultaneously willing to let a child or two die to keep them (being willing not to do something is not the same as being unwilling to do it). Suppose you refrain from sacrificing your arms without ever ceasing being willing to do so.<sup>11</sup> Now suppose that, when I face a choice just like yours, I am entirely unwilling to sacrifice my arms, and refrain from doing so. We may find it an odd implication that, while it was wrong for you not to sacrifice your arms, it was permissible for me not to sacrifice mine. It seems that if I refrained permissibly, so did you (presumably morality does not ask less of scrooges!). But again, at least according to Horton, if it is

<sup>8</sup> *Ibid.*, pp. 100–01.

<sup>9</sup> For example, F. M. Kamm recognizes and quickly rejects a Hortonesque solution to a supererogation puzzle in chapter 12 of her *Morality, Mortality, Volume II: Rights, Duties, and Status* (New York: Oxford University Press, 1996), p. 315.

<sup>10</sup> Horton, "The All or Nothing Problem," *op. cit.*, footnote 8.

<sup>11</sup> I assume it is possible to be wholeheartedly willing to do something and yet not do it. This seems to be confirmed frequently in everyday life. Consider, for example, the violent ambivalence one can experience when deciding between very different careers, or even very different restaurants.

permissible not to sacrifice your arms but wrong to save only one child, then you ought to save neither rather than save only one. That is, according to Horton, if (1) and (2), then (3). And here (3) would apply to you even though you are willing to save a child, which is what Horton finds unacceptable.

I do not claim that either of these challenges to Horton's solution constitutes a decisive reason to reject it. But they give us good reason to explore other possible solutions.

## II. CONDITIONAL OUGHTS

Whether (3) seems implausible depends on how we interpret contrastive-ought claims (claims of the form "you ought to do *A* rather than *B*").<sup>12</sup> Suppose we interpret (3) so that it captures *nothing more* than the conjunction of (1) and (2), perhaps along with a reminder that "permissible" is a better deontic status than "wrong." It is difficult to see why (3), interpreted in this minimal way, should seem implausible to us, if (1) and (2) seem plausible to us. Since we find (1) and (2) plausible but (3) implausible, presumably we do not interpret (3) in this way. In a footnote, Horton says that (3) is equivalent to:

(3\*) "[I]f you are not going to save both children, you ought to save neither."<sup>13</sup>

It is not clear that the contrastive-ought claim (3) is strictly equivalent to the conditional-ought claim (3\*), but I will from here onward interpret (3) so that it implies (3\*).

In general, conditional oughts provide action guidance for agents who are not going to perform certain acts that are available to them, including when they are not going to do what they ought to do. Even if it is settled that you ought to do *A*, there is an intelligible question about what *if you are not going to do A* you ought to do.<sup>14</sup> Suppose you can murder brutally, murder gently, or refrain from murdering altogether. Of course, you ought to refrain from murdering altogether. However, if you are not going to refrain from murdering, you ought to murder

<sup>12</sup>For an overview of relevant literature, see Justin Snedegar, "Contrastivism about Reasons and Ought," *Philosophy Compass*, x, 6 (June 2015): 379–88.

<sup>13</sup>Horton, "The All or Nothing Problem," *op. cit.*, footnote 1.

<sup>14</sup>This question arises in the debate between *actualists* and *possibilists*, though this debate tends to focus on what you ought to do now, given that, out of some set of alternatives available later, you will (not)  $\phi$ . I am focusing here on what you ought to do now, given that, out of the set of alternatives available now, you will (not)  $\phi$ . On actualism versus possibilism, see Jacob Ross, "Actualism, Possibilism, and Beyond," in Mark Timmons, ed., *Oxford Studies in Normative Ethics*, vol. 2 (Oxford: Oxford University Press, 2012), pp. 74–96; and Yishai Cohen and Travis Timmerman, "Actualism and Possibilism in Ethics," in Edward N. Zalta, ed., *Stanford Encyclopedia of Philosophy* (forthcoming).

gently. This seems like a plausible conditional-ought claim (assuming we cannot conclude that you plain old ought to murder gently, from this conditional-ought claim together with the empirical claim that you are not going to refrain from murdering). This conditional ought is known as a contrary-to-duty imperative.<sup>15</sup> There are other sorts of conditional oughts. It is permissible to say something nice, wrong to say something nasty, and permissible to say nothing at all. But, if you are not going to say something nice, you ought to say nothing at all. We can call this conditional ought a permission-excluding imperative.<sup>16</sup>

The conditional-ought claim (3\*), “if you are not going to save both children, you ought to save neither,” cannot be derived from the mere conjunction of (1) and (2). We could accept (1) and (2) but also reject (3\*). In the following section, I defend the *conditional-permission* claim that, “if you are not going to save both children, it is permissible to save only one.” (We cannot conclude that it is plain old permissible to save only one, from this conditional-permission claim together with the empirical claim that you are not going to save both children; moreover, this conditional-permission claim must be construed so as not to imply that “if it is permissible not to save both children, it is permissible to save only one.”) This conditional-permission claim does not discourage you from saving a child if you are not going to save both children but are willing to save one. (3\*) does. It is this feature of (3\*) that makes it seem implausible. It is also what makes (3) seem implausible, given that it is interpreted to imply (3\*).

### III. CONDITIONAL PERMISSIONS

Why does Horton think that two seemingly plausible claims, (1) and (2), imply a seemingly implausible conditional-ought claim, (3\*)? He appeals to the following bridge principle:

(BP) “If *A* is morally permissible and *B* is morally wrong, then we ought to do *A* rather than *B*.”<sup>17</sup>

Together with the claims that it is permissible not to sacrifice your arms and wrong to save only one child, BP implies that you ought to save neither rather than only one. That is, together with (1) and (2), BP implies (3). Since (1) and (2) are meant to imply not just (3), but also (3\*), Horton needs to appeal to:

<sup>15</sup> See Paul McNamara, “Deontic Logic,” in Edward N. Zalta, ed., *Stanford Encyclopedia of Philosophy* (Fall 2018 Edition), URL = <<https://plato.stanford.edu/archives/fall2018/entries/logic-deontic/>>.

<sup>16</sup> I owe this terminology to Joe Horton.

<sup>17</sup> Horton, “The All or Nothing Problem,” *op. cit.*, p. 96.

- (BP\*) If your alternatives are *A*, *B*, and *C*, and *A* and *C* are permissible, while *B* is wrong, then if you are not going to do *C*, you ought to do *A*.

Horton claims that BP is “intuitively correct” and that “there are countless cases that seem to verify it.” Given the way he uses BP, not to mention his claim that (3) and (3\*) are equivalent, he must think the same about BP\*.

BP\* seems correct when we consider some cases (for example, if you are not going to say something nice, you ought to say nothing at all), but it loses its intuitive appeal when we consider others. Indeed, Armsgiving is arguably such a case—the plausibility of (1) and (2), together with the implausibility of (3\*), arguably already casts significant doubt on BP\*. Consider a further case, the probative value of which should be clear by the end of this section.

*Hot Death:* You can either (I) do nothing, (II) save one hundred children by pressing a button, or (III) save these very same hundred children, and prevent a separate child from losing a foot, by pressing a different button. Pressing either button will also cause you to drop into a pit of red-hot coals, where you will die an excruciatingly painful death.

It seems that (I) is permissible.<sup>18</sup> It also seems that (II) is wrong. Given these seemingly plausible claims, BP\* implies that if you do not do (III), you ought to do (I). This conditional ought requires you *not* to save the hundred children, conditional on your not saving them along with a foot. That seems absurd. The problem lies with BP\* itself. We feed it plausible claims, and yet it implies unacceptable conditional oughts. We should, for this reason, reject it.

More modestly, *if* we find the claims that are fed in as BP\*'s antecedent to be plausible, *and if* we find BP\*'s resultant conditional-ought claims to be implausible, *then* we should locate the source of the implausibility in BP\* itself and reject it accordingly. For BP\* would, in such cases, lack independent intuitive appeal. Any initial appearance of plausibility it had, then, is likely the result of overgeneralizing from a limited range of cases. Where does BP\* err?

While *A*, *B*, and *C* are your available alternatives, we can say that, if you are not going to do *C*, *A* and *B* are your only *non-excluded* alternatives. BP\* assumes that the conditional deontic statuses of your non-excluded alternatives are a function of the deontic statuses of your available alternatives, in the following way. Since *A* is permissible, *B* is wrong, and *C* is permissible, if you are not going to do *C*, your only non-excluded alternatives are a permissible act, *A*, and a wrong act, *B*.

<sup>18</sup> If you are unsure, feel free to make your hot death more hellish still!

If your only non-excluded alternatives are a permissible act and a wrong act, you ought, given that you are not going to do the excluded act, to do the non-excluded permissible act.

This last step is where BP\* goes wrong. In Hot Death, (III) is *seriously* morally better than (II). By this I mean that (III) is morally better than (II) in a way that would make (II) wrong if there were no sufficient justification for doing it, and that there is no sufficient justification for doing it (that is, for gratuitously allowing a child to lose a foot). This is what makes (II) wrong. Even though (II) is wrong, it still seems morally better than (I), which is permissible.<sup>19</sup> While (II) is morally better than (I) in a way that would make (I) wrong if there were no sufficient justification for doing it (you let many children die), there is a sufficient justification for doing it (you avoid a hot death). So (II) is not, in the sense defined above, seriously morally better than (I). At the same time, (II) seems *substantially* morally better than (I). This is part of what makes it seem absurd that, if you are not going to do the best alternative (III), you ought *not* to do the next best thing (II). Instead it seems plausible that you would have a *conditional permission* to do (II). If you are not going to do the best thing (III), you are permitted to do the next best thing (II). Similar claims about Armsgiving also seem plausible.<sup>20</sup>

To adequately assess these conditional-permission claims and their rival conditional-ought claims (implied by BP\*), we need to get a somewhat better grip on the nature of conditional permissions.<sup>21</sup>

<sup>19</sup> For challenges to the common assumption that “if *A* is wrong and *B* is permissible, then *A* is morally worse than *B*,” see Daniel Muñoz, “Better to Do Wrong” (unpublished manuscript); and Portmore, *Opting for the Best*, *op. cit.*, section 6.4. Rejecting this assumption strengthens the case against BP\*. But even if it were true, BP\* would still seem implausible. For example, even if (II) were morally worse than (I), it would still seem plausible that (II) is permissible conditional on your not doing (III). After all, the hundred children live if you do (II), and die if you do (I).

<sup>20</sup> Jeff McMahan discusses the All or Nothing Problem in “Doing Good and Doing the Best,” in Paul Woodruff, ed., *The Ethics of Giving: Philosophers’ Perspectives on Philanthropy* (New York: Oxford University Press, 2018), pp. 78–102. Conditional permissions may enable us to accommodate the intuitions that pushed McMahan (*ibid.*, p. 100) to suggest that in Armsgiving (ii) is wrong (because another act ought to be done instead) yet not impermissible (because it would be morally better to do than some other act that is permissible). Accommodating these intuitions by appealing to conditional permissions would seem preferable to the way suggested by McMahan, since while an act can be wrong yet conditionally permissible, it is difficult to see how an act can be wrong yet not impermissible.

<sup>21</sup> My aim throughout the remainder of this section is to say just enough about conditional permissions to illuminate the broad sort of solution to the All or Nothing Problem they make possible, rather than provide anything close to a fully detailed account. The latter would address whether positing conditional permissions requires appealing to different senses or levels of “ought,” action guidance, or blameworthiness, and whether and how to differentiate between cases in which it is (at the time of

First, contrary-to-duty imperatives conditionally require you to do your “least wrong” non-excluded alternative (for example, murder gently), and permission-excluding imperatives conditionally require you to do your permissible non-excluded alternative (for example, say nothing at all). These conditional oughts provide action guidance for those who are not going to perform certain acts that are available to them. It seems there is a sense in which these agents need action guidance, and conditional oughts provide it. Conditional permissions do not similarly guide agents toward any particular non-excluded alternative. However, they can *displace* the inappropriate action guidance that otherwise would have been provided by rival conditional oughts. For example, it is inappropriate to require you not to save the hundred children given that you are not going to save them along with a foot. And (3\*) rather implausibly discourages you from saving only one child if you are not going to save both. Conditional permissions leave the agent free to choose, given what they are not going to do.

Second, being discouraged from saving only one child is compatible with not being discouraged from saving only one child *given that* you are not going to save both. Similarly, you are discouraged from murdering gently. But you are not discouraged from murdering gently given that you are not going to refrain from murdering. Of course, murdering gently is not *merely* conditionally permissible. It is conditionally what you ought to do (again, given that you are not going to do what you ought to do).<sup>22</sup> Saving only one child is merely conditionally permissible because it remains permissible to do nothing (if you could save the children at no cost to yourself, it is plausible that you ought to save both, and that if you do not save both, you ought to save only one). Murdering gently is conditionally permissible partly in virtue of your excluding what you ought to do, whereas saving only one child is conditionally permissible partly in virtue of your excluding the merely permissible alternative of saving both. But excluding what you ought to do is only a special case of the more general phenomenon of excluding alternatives in a way that is relevant to the conditional deontic statuses of your non-excluded alternatives.

---

decision) under one’s deliberative control which alternatives one excludes and cases in which it is not. As noted in footnote 14, the debate between actualists and possibilists may provide some clues. For a recent example, see Travis Timmerman and Yishai Cohen, “Moral Obligations: Actualist, Possibilist, or Hybridist?,” *Australasian Journal of Philosophy*, xciv, 4 (February 2016): 672–86.

<sup>22</sup> Just as you are permitted to do what you ought to do, you are conditionally permitted to do what you conditionally ought to do.



Third, as suggested parenthetically in section II, not all interpretations of conditional oughts leave room for the sort of conditional permissions needed to solve the All or Nothing Problem. We need a conditional-permission claim that is incompatible with (3\*) but compatible with the thought that, since it is permissible to save neither child, wrong to save only one, and permissible to save both, you ought to save both children or save neither. Yet on a wide-scope interpretation of conditional oughts, (3\*) is equivalent to “you ought to (if not save both children, save neither),” which is equivalent to “you ought to (save both children or save neither),” which is equivalent to “you ought to save both children or save neither.” Fortunately, there are alternatives to this wide-scope interpretation. On a standard dyadic interpretation, “if you do (not do) *A*, you ought to do an alternative in set *S*” is true if and only if *S* is the set of all the morally best alternatives in the (non-singleton) superset of alternatives not excluded by (not) doing *A*.<sup>23</sup> We can modify this interpretation of conditional oughts so that *S* includes all the morally best alternatives that are not excluded *and* all the unconditionally permissible alternatives that are not excluded (if there are any). That way, if you do not save both children, you ought to save only one child (the morally best non-excluded alternative) or save neither (the unconditionally permissible non-excluded alternative). You are thus permitted to save only one, conditional on not saving both.<sup>24</sup>

Fourth, the interpretation of conditional oughts just suggested also implies that, if you save at least one child, you ought to save both (the morally best and unconditionally permissible non-excluded alternative). This result is plausible. We cannot conclude from this conditional-ought claim, together with the empirical claim that you are going to save at least one child, that you plain old ought to save both children. This conditional-ought claim is compatible with (1), whether or not you are willing to sacrifice your arms to save a child.

Fifth, our account of conditional permissions need not preclude the availability of excluded alternatives from affecting the conditional deontic statuses of non-excluded alternatives. In Armsgiving, you can either (i) do nothing, (ii) save one child by sacrificing your arms, or (iii) save both children by sacrificing your arms. Suppose we accept that there are conditional permissions and claim that, if you are not

<sup>23</sup>For a recent discussion of wide-scope, dyadic, and restrictor interpretations of conditional oughts, see Juan Comesaña, “Normative Requirements and Contrary-to-Duty Obligations,” this JOURNAL, CXII, 11 (November 2015): 600–26.

<sup>24</sup>I am grateful to Krister Bykvist, Kevin Scharp, and Justin Snedegar for useful discussions of the relation between various interpretations of conditional oughts and the account of conditional permissions defended here.

going to do (iii), you are permitted to do (ii). We may also claim, as seems true, that if (iii) were unavailable, (ii) would be permissible. It would be a mistake to move from the latter claim to the former claim by treating excluded alternatives as if they were unavailable. For one thing, this would conflate conditional permissibility and plain old permissibility. For another, it would make the account of conditional permissions unnecessarily restrictive. As an illustration of the latter, consider a version of Armsgiving in which your alternatives are (i\*) do nothing, (ii\*) save a child's left arm by sacrificing both your arms, or (iii\*) save this very same child's left arm, and save another child's life, by sacrificing both your arms. Many claim that there is not only optionality to favor yourself over strangers, but also optionality to favor strangers over yourself (where "favoring" in either case departs from what would be impartially best).<sup>25</sup> They might believe that, if (i\*) and (ii\*) were your only alternatives, it would be permissible for you to do either. Nonetheless the availability of (iii\*) arguably makes (ii\*) so disrespectful to the child whose life you could save that it makes (ii\*) seriously morally worse than (i\*).<sup>26</sup> Suppose this is correct. It then seems that if your only *non-excluded* alternatives are (i\*) and (ii\*), you conditionally ought to do (i\*). That is, you *lack* permission to do (ii\*), conditional on not doing (iii\*). If instead we treated excluded alternatives as if they were unavailable, we would then claim that, since (iii\*) is unavailable, (ii\*) is not disrespectful, and is therefore conditionally permissible. Even if this claim were true, our account of conditional permissions need not be so restrictive as to rule out the opposite claim by generally precluding the availability of excluded alternatives from affecting the conditional deontic statuses of non-excluded alternatives.

Sixth, perhaps in the original version of Armsgiving (ii) is disrespectful to the child whose life you do not save. But it seems this would not make (ii) morally worse than (i). Even if disrespectful, (ii) remains a rather heroic life-saving act. It seems to be the fact that (iii) is seriously morally better than (ii) that makes (ii) wrong.<sup>27</sup> Hot Death is a clearer counterexample to BP\* than Armsgiving partly because it is clearer that the only fact sufficient to make (II) wrong is that (II) is

<sup>25</sup> For example, Thomas Hurka and Esther Shubert, "Permissions to Do Less Than the Best: A Moving Band," in Timmons, ed., *Oxford Studies in Normative Ethics*, vol. 2, *op. cit.*, pp. 1–27.

<sup>26</sup> For discussion of the role of disrespect in Armsgiving-type cases, see Christian Barry and Seth Lazar, "Beyond the First Call of Duty: Supererogation and Optimisation" (unpublished manuscript).

<sup>27</sup> This is compatible with claiming that the fact that (ii) is disrespectful further contributes to (ii)'s wrongness.

seriously morally worse than (III) than it is that the only fact sufficient to make (ii) wrong is that (ii) is seriously morally worse than (iii). In addition, insofar as (ii) is morally better than (i), (II) is morally better than (I) to a greater degree. (II) is less disrespectful, more heroic, and saves many more lives than (ii). Requiring you not to do the next best thing, conditional on your not doing what is best, is even more implausible in Hot Death than it is in Armsgiving.

Finally, Armsgiving and Hot Death are not the only cases that make trouble for BP\*. Indeed, many cases in which it is plausible that *expansion consistency* fails seem to provide fairly straightforward recipes for generating seeming counterexamples to BP\*. Consider the version of expansion consistency according to which if *A* and *B* are permissible when the available alternatives are *A* and *B*, then *A* is permissible if and only if *B* is permissible when the available alternatives are *A*, *B*, and *C*.<sup>28</sup> This condition fails in Armsgiving, since if your only alternatives are doing nothing and sacrificing your arms to save one child, each is permissible, but if you also have the alternative of sacrificing your arms to save both, saving only one is wrong while doing nothing remains permissible. It is plausible that it fails in various other cases too.<sup>29</sup>

#### IV. A NEXT BEST MORALITY WITH OPTIONALITY

I believe that the foregoing discussion provides insight into what the All or Nothing Problem was all along. The problem is not simply how to combine the non-consequentialist optionality behind claims like (1) with the view behind claims like (2) that it can be wrong to do good but suboptimal acts. The problem is how to satisfactorily combine both these elements together with conditional oughts—if you are not going to save both children, what ought you to do?

Horton's view of conditional oughts, according to which BP\* is true, leads us from the plausible claims (1) and (2) to the implausible conditional-ought claim (3\*), that, if you are not going to save both children, you ought to save neither. Both Horton and the act consequentialist avoid (3\*) by rejecting (1), at least in cases in which you

<sup>28</sup> See Amartya Sen's "Property  $\beta$ " in "Quasi-transitivity, Rational Choice and Collective Decisions," *Review of Economic Studies*, xxxvi, 3 (July 1969): 381–93, at p. 384.

<sup>29</sup> Kamm's supererogation puzzle is another case in which it is plausible that expansion consistency fails (*Morality, Mortality, Volume II, op. cit.*, chapter 12). So are cases involving incommensurability or parity; see Ruth Chang, "Value Incomparability and Incommensurability," in Iwao Hirose and Jonas Olson, eds., *Oxford Handbook of Value Theory* (Oxford: Oxford University Press, 2015), pp. 205–24. For further discussion, see Tina Rulli and Alex Worsnip, "IIA, Rationality, and the Individuation of Options," *Philosophical Studies*, clxxiii, 1 (January 2016): 205–21; and Daniel Muñoz, "Supererogation and Rational Choice: Incommensurability, Intransitivity, Independence" (unpublished manuscript).

are willing to sacrifice your arms to save a child. This enables them to say, consistently with BP\*, that if you are not going to save both children, you ought to do the next best thing, which is to save only one. On their view, this is the morally best (or least wrong) non-excluded alternative.

I have argued that we should reject BP\*. This principle lacks independent intuitive appeal in the very sorts of cases that undergird the All or Nothing Problem. Rejecting BP\* enables us to accept (1) and (2), reject (3\*), and instead accept the conditional-permission claim that, if you are not going to save both children, you are permitted to do the next best thing, that is, to save only one. Here is an equivalent summary of my view. You ought to save both or save neither, but if you are not going to save both, you ought to save one or save neither.<sup>30</sup>

Now suppose your alternatives are (i) do nothing, (ii) save one child by sacrificing your arms, (iii) save this very child and another by sacrificing your arms, or (iv) save both these very same children and a third by sacrificing your arms. On a plausible extension of what I have argued here, you ought to save all three or save none, but if you are not going to save all three, you ought to save two or save none, and so on. It would be implausible to claim that, because you have excluded saving all three, you are conditionally permitted to do any of these non-excluded alternatives. Of these non-excluded alternatives that are wrong, you are conditionally permitted to do only the one that is morally best (or least wrong). At the same time, you are permitted to save none.

In general, when you are permitted not to do any good in virtue of the large sacrifice needed of you to do so, and the act that does the most good requires no more of you than (and is seriously morally better than) the act that does the next most good, and the act that does the next most good requires no more of you than (and is seriously morally better than) the act that does the *next* next most good, and so on, you ought to do the most good or none at all, but if you do not do the most good, you ought to do the next most good or none at all, and so on. This appears to me an attractive way of combining optionality and conditional oughts together with the view that it can be wrong to do good but suboptimal acts.

<sup>30</sup> Ralf Bader offers a very concise and technically precise statement of this view in footnote 28 of his "Agent-Relative Prerogatives and Suboptimal Beneficence," in Mark Timmons, ed., *Oxford Studies in Normative Ethics*, vol. 9 (Oxford: Oxford University Press, forthcoming). We arrived at the view independently, though after seeing Bader's footnote I benefited from conversations with him.

As this paper comes to an end, it is worth remembering that the All or Nothing Problem and the next best morality that can solve it are not confined to cases in which next best acts are dominated. Consider a version of Armsgiving in which six children face a deadly threat, and your alternatives are (i\*\*) do nothing, (ii\*\*) save one child by sacrificing your arms, or (iii\*\*) save the other five by sacrificing your arms. This version of Armsgiving differs from the original in that the next best act (ii\*\*) is not dominated by the best act (iii\*\*). Though (iii\*\*) is overall better than (ii\*\*), (ii\*\*) is in a way better than (iii\*\*). If you do (ii\*\*), you save a particular child who would have died if you had done (iii\*\*) instead. The claim that (ii\*\*) is wrong even if (i\*\*) is permissible is admittedly less intuitive than the claim that in the original version of Armsgiving (ii) is wrong even if (i) is permissible. Still we should accept both these claims. It is wrong to save the one rather than the five when neither act requires any sacrifice and all other things are equal.<sup>31</sup> And it seems that if it is wrong to save the one rather than the five when neither act requires any sacrifice and all other things are equal, then it is also wrong to save the one rather than the five when both require the same sacrifice and all other things are equal.<sup>32</sup> Even though (ii\*\*) is not dominated by (iii\*\*), it is still the case that (ii\*\*) is seriously morally worse than (iii\*\*). Yet (ii\*\*) is morally better than (i\*\*). It would again be implausible to claim that, if you are not going to do (iii\*\*), you ought not to do the next best thing (ii\*\*). As before, you are conditionally permitted to do (ii\*\*) if you are not going to do (iii\*\*). And, as before, you remain permitted to do (i\*\*).

One final remark. Horton and I have independently argued that the claim that it is wrong to make a sacrifice in a way that saves a few nearby strangers rather than many others has important implications for charitable giving.<sup>33</sup> We have argued that there are many cases in which it is wrong to engage in suboptimal charitable giving (even if it is

<sup>31</sup> While a famous few reject this claim, the overwhelming majority of people accept it. For survey data, see Mark Kelman and Tamar Kreps, "Playing with Trolleys: Intuitions about the Permissibility of Aggregation," *Journal of Empirical Legal Studies*, xi, 2 (June 2014): 197–226. For relevant literature, and an explanation of why you ought to save the many and not the few, see Tom Dougherty, "Rational Numbers: A Non-consequentialist Explanation of Why You Should Save the Many and Not the Few," *Philosophical Quarterly*, LXIII, 252 (July 2013): 413–27.

<sup>32</sup> Perhaps we enjoy prerogatives to use what is rightfully ours as we please, independently of any costs to us. I believe that, if we have such prerogatives, they are limited. They would seem insufficient to justify your using what is rightfully yours to save one person rather than many others when these acts are equally costly to you and all other things are equal. For discussion of property-based prerogatives, see Fiona Woolard, *Doing and Allowing Harm* (Oxford: Oxford University Press, 2015), chapters 7 and 8.

<sup>33</sup> See my "Whether and Where to Give," *op. cit.*; and Horton, "The All or Nothing Problem," *op. cit.*, section iv.

permissible to do still less good by simply not giving). Our arguments appeal to the premise that it is wrong to do much less good rather than much more, if doing more is no costlier to you, and all other things are equal. While this premise is compatible with non-consequentialism, it is controversial among more thoroughgoing non-consequentialists.<sup>34</sup> Indeed, it is unnecessarily strong for our purposes. But for all that, it would seem to us incredible that, though it is wrong to sacrifice your arms in a way that saves a few nearby strangers rather than many others, it is not wrong to give a sum of money to charity in a way that saves a few distant strangers rather than many others.<sup>35</sup>

THERON PUMMER

University of St Andrews

<sup>34</sup>Thomas Sinclair, "Are We Conditionally Obligated to Be Effective Altruists?," *Philosophy and Public Affairs*, XLVI, 1 (Winter 2018): 36–59.

<sup>35</sup>I am grateful to Joe Horton for discussion of this last paragraph.