

12-2008

On the Concept of a Morally Relevant Harm

David Lefkowitz

University of Richmond, dlefkowi@richmond.edu

Follow this and additional works at: <http://scholarship.richmond.edu/philosophy-faculty-publications>

 Part of the [Ethics and Political Philosophy Commons](#)

This is a pre-publication author manuscript of the final, published article.

Recommended Citation

Lefkowitz, David, "On the Concept of a Morally Relevant Harm" (2008). *Philosophy Faculty Publications*. 61.
<http://scholarship.richmond.edu/philosophy-faculty-publications/61>

This Post-print Article is brought to you for free and open access by the Philosophy at UR Scholarship Repository. It has been accepted for inclusion in Philosophy Faculty Publications by an authorized administrator of UR Scholarship Repository. For more information, please contact scholarshiprepository@richmond.edu.

Abstract: On the Concept of a Morally Relevant Harm

In this paper I explicate and defend the concept of a morally relevant harm. This concept figures prominently in common-sense and contractualist moral reasoning concerning cases where an agent can prevent harm to members of a large group or a small one, but not both. When the two harms to which members of these groups are exposed are morally relevant to one another, an agent is permitted (or perhaps required) to take into account the number of people he can save. When the harms are irrelevant, an agent should not even consider preventing the lesser harm, regardless of how many people will suffer it.

I argue for what I label the orbital conception of morally relevant harm, according to which harms that fall within the “orbit” of a given harm are relevant to it, while all other harms are not. In addition, I contend that the possibility of preventing a harm provides both a first-order reason to prevent that harm, and a second-order reason not to consider preventing irrelevant harms. I then demonstrate how this understanding of the concept of a morally relevant harm avoids two objections raised by Alastair Norcross: first, identifying a point along a continuous scale of harms at which the divide between relevant and irrelevant harms occurs, and second, the entailment that the mere possibility of being able to prevent harm that one is morally forbidden from preventing can determine which of two other actions morality requires.

Keywords: morally relevant harm; contractualism; consequentialism; Scanlon; Norcross; exclusionary reason.

On the Concept of a Morally Relevant Harm

According to common-sense morality, given the choice between preventing one person's death or five people's death, an agent has a moral duty to do the latter. But given the choice between preventing one person's death or any large but finite number of slight, temporary, headaches, an agent has a moral duty to do the former. That is, there is no number of headaches such that the possibility of preventing them can even make it permissible, let alone a duty, to prevent those headaches rather than one death.

Moreover, this is not because as a practical matter there could never be enough creatures who could suffer slight, temporary, headaches such that the harm they would suffer would provide a reason to prevent the headaches that would outweigh or defeat the reason provided by the harm one person threatened with death would suffer. Rather, the prevention of a slight, temporary, headache is simply not the kind of consideration an agent ought to take into account in circumstances where he can prevent a death, and the number of people who will suffer headaches does nothing to change this.

Yet it does not seem that common-sense morality always bars consideration of how many people will suffer a certain kind of harm given the possibility of preventing someone from suffering a greater harm. For instance, given the possibility of preventing a large number of people from suffering total and permanent paralysis, or preventing one person's death, common-sense morality directs us to prevent the former, rather than the latter.¹ It seems, then, that at least in conflict cases involving (only) the prevention of harm, common-sense morality instructs us to take the possibility of preventing lesser harms into account in some cases, but not others. Or, as T. M. Scanlon writes in *What We Owe to Each Other*,

if one harm, though not as serious as another, is nonetheless serious enough to be morally “relevant” to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even “relevant to,” some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm.²

My purpose in this paper is to explicate and defend the concept of a morally relevant harm (or the relation ‘is a morally relevant harm to’), a task I pursue primarily by way of responding to criticisms of it recently advanced by Alastair Norcross.³

Norcross criticizes the concept of morally relevant harms because it denies the transitivity of harms. In addition to decrying the sheer implausibility of such a denial, Norcross identifies two problems for the intransitive notion of morally relevant harms. The first of these is the difficulty in identifying the point along a continuous scale of harms at which a break occurs, particularly since any such break appears to entail that two harms that differ hardly at all will fail to be morally relevant to one another. The second is the “strange” and “unpalatable” conclusion that in a certain kind of three-option case, the mere physical possibility of being able to prevent harm *that one is morally forbidden from preventing* can determine which of two other actions one is morally required to take.

To meet Norcross's first objection, I defend what I label an orbital conception of morally relevant harm, according to which the harms that fall within the "orbit" of a given harm are relevant to it, while all other harms are not. In addition, I construe the reason for action provided by the orbital conception of morally relevant harm as consisting of both a first-order reason to prevent that harm, and a second-order reason not to consider the possibility of preventing irrelevant harms, regardless of how many people stand to suffer it. Such a construal prevents a modified argument for the transitivity of harms (at least for purposes of an agent's deliberation), and also plays a key role in my response to Norcross's second objection. Specifically, the second-order reason not to consider the possibility of preventing irrelevant harms transforms Norcross's problematic three-option case into two two-option cases, which can both be easily addressed using the principle for choice in conflict cases that Scanlon defends.

I offer two justifications for thinking that the option of preventing a given harm provides a second-order reason not to consider the possibility of preventing certain other harms. First, I contend that such a conception of the reason provided by the option of preventing a harm cannot be reasonably rejected from any generic standpoint in Norcross's allegedly problematic three-option case. Such an argument obviously depends on the justifiability of Scanlon's contractualism. Yet Scanlon's remarks on the concept of a morally relevant harm are separable from his moral theory; moreover, as indicated above, something like the concept of a morally relevant harm appears to be part of common-sense morality, which we should not assume is merely a rough form of contractualism. Therefore, I sketch a non-contractualist account of why agents ought to treat certain harms as providing a second-order reason not to consider the possibility of

preventing certain other harms, one that focuses on the way in which considering how many people will suffer a lesser harm trivializes the loss to the person at risk of the greater harm.

I

Suppose we have a continuous scale of harm stretching from A to Z, with each harm differing from the one before it only as much as is necessary for it to count as a distinct harm. If harms D and E are morally relevant to one another, then even though D is a worse harm than E, in some cases an agent will be morally permitted, and perhaps even required, to prevent a large number of people from suffering harm E rather than a small number of people from suffering harm D, at least when he can only help one party or the other. Or at least this is what Scanlon suggests, and for present purposes I shall assume the truth of his claim.⁴

It seems highly plausible that any harm will at least be morally relevant to the harms just before and just after it along the scale. If so, however, then transitivity entails that harm A will be relevant to harm Z, since A is relevant to B, B to C, and so on. But this entails a conclusion that both common-sense morality and Scanlon reject, namely that when faced with a choice between preventing one person's death and preventing a number of people from experiencing headaches, we ought to take into account how many people will suffer the harm in question (or worse yet, aggregate the harms in question), with the in principle possibility that we ought to prevent the headaches rather than the death. As the earlier quotation suggests, Scanlon objects not to the judgment that *n* number of headaches suffice to outweigh one life, but rather to the very consideration of headaches in the moral deliberation of an agent who can prevent one person's death.

Perhaps someone might argue that preventing others' headaches, even when I can do so at little cost to myself, falls outside what we owe to each other; that is, it qualifies as supererogatory, rather than as a moral duty. Even if this is true in the case of headaches, it seems possible to identify cases where I have a duty to prevent a harm, at least if I can do so at little cost to myself, and yet given the possibility of preventing some other harm, I ought to treat the first harm as irrelevant. For example, it seems to me that I have a duty to prevent one person suffering a broken arm, if I can do so at little cost to myself. Yet it also seems that given the possibility of preventing one person's death, I ought not to even consider the possibility of preventing broken arms, regardless of how many broken arms I might prevent.

Apparently, then, anyone attracted to the notion of a morally relevant harm will have to deny the transitivity of harms if he wishes to avoid the in principle comparison of lives and headaches. But where along the continuous scale does the break in transitivity occur? Consider one possible point: between L and M. By stipulation, there is almost no difference between these two harms. Yet if M is not morally relevant to L, then an agent faced with a choice between preventing one person from suffering L, and two, three, or a million people from suffering M, is morally required to choose the first course of action. Such a conclusion is surely false. Norcross suggests that this will be so even for the intuitively most attractive point for locating an intransitivity, namely death. As he writes, "can anyone who really considers the matter seriously honestly claim to believe that it is worse that one person die than that the entire sentient population of the universe be severely mutilated?"⁵

This illustration of the difficulty in identifying a point along a continuous scale of harms at which the relationship of transitivity does not hold rests on the assumption that the notion of moral relevance establishes two or more categories of harm, with particular harms along the scale falling in only one category. Call this the **categorical** conception of morally relevant harms. To take the example mentioned in the previous paragraph, there might be two categories of harms, 1 and 2, with harms A through L falling within category 1, and harms M through Z falling in category 2. All of the harms within a given category are relevant to one another, and the relationship of transitivity holds between them. However, no harm in category 1 is relevant to any harm in category 2. Given the chance to prevent a harm in category 1, an agent should not even consider the possibility of preventing harms in category 2. That is, he ought not to recognize any transitivity between harms in category 1 and harms in category 2. The obvious problem with this conception of morally relevant harms is that it will inevitably be arbitrary where exactly along the scale of continuous harms we establish a boundary between different categories of harm. Why is it, given that there is so little difference between L and M, that L is a category 1 harm, while M is a category 2 harm? Indeed, the difference between L and M is significantly less than the difference between M and Z, and yet in some cases we may (or even must) take the numbers into account when we can prevent either M or Z (but not both), while we are forbidden from doing so when we can prevent either L or M.⁶

Fortunately, denying transitivity across the entire scale of harms does not require that we adopt the categorical conception of relevant harms. Recall Scanlon's suggestion that for any harm, those harms that are roughly akin, though not equivalent, to it will count as morally relevant. Modeled very simply, we can say that harm A is relevant to B

and C, but not D, B is relevant to harms A, C, and D, but not E, and so on. Call this the **orbital** conception of a morally relevant harm; any harm that is within the orbit of a given harm is relevant to it, while any harm that is not within the orbit of a given harm is not relevant to it. The orbital conception of relevance entails that harm Z is not relevant to harm A, since the former does not fall within the orbit of the latter. Yet it avoids treating two harms that differ very little from one another, such as L and M, as not morally relevant to one another, at least on the plausible assumption that, for any harm, those harms that differ from it only a little will fall within its orbit.

The orbital notion of morally relevant harms may still appear to entail the transitivity of A and Z. For even though Z is not within the orbit of harm A, it is within the orbit of harm X, which is within the orbit of harm V, and so on until we reach harm C, which is within the orbit of harm A. Given that transitivity holds amongst relevant harms, it appears to follow that harms A and Z can be compared. However, this conclusion rests on a failure to properly appreciate the nature of the reason provided by the orbital conception of a morally relevant harm.

Like the categorical conception, the orbital conception of a morally relevant harm reflects the boundary separating those harms that an agent ought to treat as comparable in their moral seriousness from those he ought not to treat as comparable. The fact (if it is one) of such a boundary cannot be captured by a conception of practical reason that treats all reasons as merely weighty considerations, with practical reasoning consisting of the weighing of reasons for and against a particular action. Such a conception of reasons and practical reasoning treats all harms as comparable. Rather, we must employ a two-level conception of practical reasoning, according to which harms (as well as other

considerations) can provide either a first-order reason for action, or a second-order reason to refrain from considering certain first-order reasons for action, or both.⁷ The concept of a second-order reason that excludes from an agent's deliberation certain first-order reasons for action captures the idea of a boundary delimiting harms comparable in their moral seriousness from those that are not.⁸ So for example, the option of preventing harm A excludes from an agent's deliberation the option of preventing any harms that are not relevant to A, including Z. Or, to make the example slightly more concrete, the option of preventing a death excludes from an agent's deliberation the option of preventing any harms that are not relevant to death, including headaches.

Scanlon could well be expected to avail himself of this understanding of the way in which the concept of a morally relevant harm functions in an agent's practical reason, since he clearly endorses second-order exclusionary reasons as part of normal human practical reasoning. Apropos of the issue under discussion, Scanlon notes in his discussion of the structure of reasons that

one consideration, C, [may] be a reason for taking another consideration, D, not to be relevant to my decision whether or not to pursue a certain line of action. Often our judgment that a certain consideration is a reason builds in a recognition of restrictions of this kind at the outset: D may be taken to be a reason for acting only as long as considerations like C are not present.⁹

Moreover, a sophisticated consequentialist can consistently adopt a two-level conception of practical reasoning, though of course the justification he gives for employing it will differ from that a contractualist offers.¹⁰ Thus a consequentialist such as Norcross cannot

object to the understanding of a morally relevant harm sketched here that it begs the question against consequentialism by assuming a non-consequentialist, or at least deontological, theory of practical reason.

How do we know whether some lesser harm falls within, or outside of, the orbit of some other harm? Such a determination is ultimately a matter of judgment, but it seems that we can at least say the following in response to this question. First, for many harms (and perhaps even for any harm), there is at least one harm, and likely more than one, that falls within their orbits, and at least one harm, and likely more than one, that clearly does not fall within their orbits. If so, then in at least some possible cases in which a person can prevent one of two harms, but not both, it will be clear that one of the harms is not relevant to the other, and so the agent ought not to give any consideration to preventing that harm in his deliberation. Second, cases may arise in which reasonable people disagree as to whether a given harm that it is possible to prevent falls within the orbit of some other harm that it is also possible to prevent. Disagreement is reasonable when it results from agents drawing different conclusions under what Rawls labels the burdens of judgment.¹¹ In such cases, a particular sort of procedural response may be morally required. For instance, it may be that a certain kind of political process should be employed to settle, at least for action guiding purposes, whether the lesser harm falls within the orbit of the greater one.¹² Alternatively, or perhaps in addition, it may be that as long as a person's judgment regarding the relevance of one harm to another is reasonable and sincere, then no one has a claim (right) against his acting on that judgment, even if that judgment is not actually correct.¹³ As both of these comments suggest, reasonable disagreement over the relevance of two harms need not undermine

the attractiveness of the orbital conception of a relevant reason as part of a plausible account of practical rationality.¹⁴

Third, it is possible that on the orbital conception, the concept of a relevant harm will turn out to be vague, or even somewhat indeterminate. If so, then in some cases there may be no correct answer to the question ‘does harm A fall within or outside of the orbit of harm B?’¹⁵ Perhaps some will think this possibility shows that the orbital conception of a relevant harm is mistaken, but I see no reason to assume that there must always be a determinate answer to the aforementioned question. Moreover, I suspect that any indeterminacy that might arise will do so only at the very edge of any given harm’s orbit, and so I do not think we are likely to frequently encounter cases in which it is indeterminate whether a given harm is relevant to another. Finally, just as it may be morally necessary to implement particular procedures in order to settle reasonable disagreements over the relevance of two harms, so too we may need such procedures in order to address occasional cases of indeterminacy.

In sum, according to the orbital conception of morally relevant harm, the option of preventing a given harm provides an agent with both a first-order reason to prevent that harm, and a second-order reason not to consider as part of his deliberation the option of preventing irrelevant harms. For the purposes of deliberation, then, the orbital notion of morally relevant harm does entail the denial of transitivity between harms A and Z. Given the possibility of preventing harm A, an agent who reasons properly ought not to even consider the possibility of preventing harm Z, regardless of how many people will suffer that harm, as doing so constitutes a failure to properly appreciate the reason provided by the possibility of preventing harm A. It may be worth noting that, as a

matter of metaphysics, the agent need not deny the transitivity of harms A and Z; he must merely deny its relevance to his practical reasoning. Perhaps no satisfactory justification can be offered for such a denial, but it seems a less daunting task than arguing for the metaphysical claim, as Norcross at times seems to suggest Scanlon, or any other defender of the concept of morally relevant harm, must do.¹⁶

II

Having addressed Norcross's objection regarding where the break in the spectrum of harms occurs, I turn now to the second problem he identifies as arising out of the denial of the transitivity of harms. His objection focuses on a special kind of case in which an agent is confronted with three or more parties for whom he can prevent some harm, but where preventing harm to one of these parties entails not preventing harm to the others. Suppose that an agent must choose between the following three options: (1) preventing one person from suffering harm A; (2) preventing five people from suffering harm C; and (3) preventing 50 people from suffering harm E. Suppose further that harm A is relevant to harm C, but not to harm E; harm C is relevant to both harms A and E, and harm E is relevant to harm C, but not to harm A. Let us assume as well that in cases where an agent must choose between relevant harms, he is morally required to prevent the larger number from suffering the relevant harm to which they are exposed.¹⁷ It appears to follow, then, that between options 1 and 2, the agent ought to choose option 2, or in other words prevent five people from suffering harm C rather than one person from suffering harm A. But between options 2 and 3, the agent ought to choose option 3, while between option 1 and 3 the agent ought to choose option 1. Thus we have the

following ordering of options an agent has a duty to choose: $1 < 2 < 3 < 1$, where the symbol “<” stands for “loses to.”

One obvious problem with this conclusion is that it presents the agent with a circular ordering.¹⁸ For any one of the options open to the agent, there is some other possible option that the agent is morally required to choose instead. The orbital conception of morally relevant harms allows us to avoid this circular ordering, however. Among the three options open to the agent, A is the worst harm he can prevent. Therefore, according to the orbital conception of relevant harms, the agent should only consider in his deliberation the possibility of preventing harms that are relevant to A. In the case at hand, that means that the agent may not consider the possibility of preventing E, since this is not a harm that is relevant to A. In the terms of the two-level theory of practical rationality sketched earlier, the option of preventing harm A provides a second-order reason for the agent to exclude from his deliberation the option of preventing harm E, as well as a first-order reason to prevent harm A. Thus the special three-option case that appears to produce a circular ordering, and so no answer to the question “what is an agent who confronts such a situation morally required to do?” is reduced to a two-option case in which the demands of morality are clear. Since harm C is relevant to harm A, the agent ought to take into account how many people will suffer the respective harms, and so prevent the five people from suffering harm C rather than one person from suffering harm A.

Note that the problem posed by a three-option case of the kind under consideration is not merely a matter of an agent confronting three options, but rather having three (or more) options with the kind of relevance ordering described above. So

for example, there is no difficulty with a three-option case in which an agent can prevent one person from suffering harm A, two from suffering harm B, or three from suffering harm C. Since harms B and C are both relevant to A, the agent ought to treat them as equally morally serious, and so he is required to save the greatest number.

However, in certain circumstances this account of why the apparently problematic three-option case really constitutes an easily resolvable two-option case can give rise to what Norcross labels “strange” and “unpalatable consequences” that any plausible moral theory should seek to avoid.¹⁹ The following example can be used to make his case. A moral agent, confronted with options 1, 2, and 3, chooses option 2, and initiates a course of conduct that will result in his preventing five people from suffering harm C. However, soon after he adopts this course of action, it becomes impossible for the agent to pursue option 1; that is, for him to prevent one person from suffering harm A. For example, shortly after choosing to pursue option 2, the gas tank on the agent’s car develops a small leak, so that the agent will not have enough fuel to successfully pursue option 1, though he will have enough to pursue either option 2 or 3. Mere chance has changed the agent’s situation from one in which he could choose options 1, 2, or 3, into a situation in which he can choose only between options 2 or 3. But whereas in the former situation the agent was required to choose option 2, in the latter situation the agent is required to choose option 3.

That mere chance might change an agent’s moral obligations is neither strange nor unpalatable, of course; there is no reason to reject a moral theory that requires a person to change course and save a smaller number of people from death when, by chance, it becomes impossible for him to save a larger number from the same fate. What

does appear to be strange, and at least in Norcross's view, unpalatable, is that the elimination from the realm of possibility of an action that an agent *was morally forbidden from pursuing* should determine what course of action the agent is morally required to do. That is, whereas the rescuer had a duty to prevent five people from suffering harm C, because it became impossible for him to prevent one person from suffering harm A, *which he was forbidden from preventing anyway given the possibility of preventing more people from suffering a relevant harm*, the agent now has a duty to prevent fifty people from suffering harm E. As Norcross puts it, "when one of the forbidden alternatives by chance becomes unavailable, ...the other forbidden alternative becomes obligatory, and the previously obligatory alternative becomes forbidden."²⁰

In response, a contractualist such as Scanlon can ask whether the intuitive judgment that the concept of a relevant harm entails strange and unpalatable conclusions in such cases provides grounds on which to reasonably reject that concept (or perhaps better, those principles that employ it). If not, then it is unclear whether the contractualist should be troubled by these entailments. Moreover, the very process of demonstrating why none of the agents in a special three-option case can reasonably reject the concept of a relevant harm may help to lessen the sense that its entailments in such a case are strange and unpalatable.

The key question, it seems, is whether the concept of a relevant harm can be reasonably rejected from the generic standpoint of an agent who stands to suffer harm C, for it is such an agent who seems to lose out when the gas tank in the rescuer's truck springs a leak. This agent cannot reject the orbital conception of relevant harm, since this is necessary to deny that one is obligated to prevent a large but finite number of

headaches rather than one death, without having to endorse the view that two harms that differ hardly at all are not relevant to one another. Nor can this agent reject the exclusionary role played by option 1 in the rescuer's deliberation, since this element is necessary both to prevent a modified transitivity of all harms argument, and so the tradeoff between a life and n headaches, and more importantly for the example in question, to block a circular ordering of moral obligations to which the rescuer is subject. The exclusionary element provided by the option of preventing harm A to one person entails that three-option cases with circular orderings that would otherwise either pose a moral dilemma or make any choice of what to do irrational (if there is a difference between the two), are instead two option cases in which, by assumption, no agent can reasonably reject the rescuer preventing harm to the greater number.

If the above argument is correct, then it is not reasonable to reject the concept of a relevant harm, understood along the lines I have described herein. But for the contractualist, the primary (and perhaps even sole) reason for concluding that a moral principle is "unpalatable" is if it could be reasonably rejected. As for Norcross's claim that the notion of a relevant reason leads to a "strange," meaning counter-intuitive, entailment in special three-option cases, this entailment seems no more strange than some of those a Utilitarian like Norcross accepts, not least that there is a large but finite number of slight, temporary, headaches such that we ought to prevent them rather than one person's death. Moreover, three-option cases like the one under consideration are likely to be extremely rare in practice. If so, then if a contractualist theory that includes the notion of a morally relevant harm captures a wide range of our considered judgments, its

inability to do so in this one case should count very little, if at all, against the acceptance of such a theory.

As I noted in the introduction, Scanlon's remarks concerning the concept of a morally relevant harm are separable from his contractualist moral theory.²¹ In the remainder of this paper, therefore, I attempt to sketch a non-contractualist justification for the orbital conception of morally relevant harm, one open to those who would defend common-sense morality (or some other non-consequentialist moral theory) on other grounds. While it may be that a contractualist can appeal to this argument as well, one need not accept contractualism in order to do so.

Much, if not all, of the sense that the concept of a morally relevant harm leads to a strange and unpalatable result in Norcross's three-option case stems from a focus on the outcome – that is, on who is saved. Given such a focus, it seems quite reasonable to wonder why the possibility or impossibility of saving a person you are morally forbidden from saving anyway should determine which of two other groups of people you ought to rescue. If instead we shift our focus to how an agent ought to deliberate about what to do (or as the contractualist would say, given that he wishes to act only in ways that no one can reasonably reject), the fact that option 1 - the possibility of preventing harm A to one person – factors in our deliberation, even though ultimately we ought not to pursue it, may not seem strange after all.

In ignoring the option of preventing harm E, regardless of how many people will suffer it, we recognize the significance of harm A to the agent who is about to incur it. In contrast, to consider how many people will suffer harm E as part of our process of determining whom we ought to save from harm trivializes the loss to the person who

stands to suffer harm A. An unwillingness to consider certain possible courses of action, given the possibility of preventing harm A, is simply part of what it is to recognize, and properly value, the loss an agent will incur if he suffers that harm. That properly valuing something sometimes involves an unwillingness to consider certain possible courses of action ought not to strike us as strange. Such an attitude is constitutive of friendship, for instance, which by its nature rules out of bounds certain calculations of self-interest on which agents are otherwise morally permitted to act. But neither does properly valuing something necessarily entail an unwillingness to consider *any* other possible course of action. Thus in Norcross's three-option case, when we fail to prevent one person from suffering harm A, that is only because we also recognize (1) the significance of harm C to the five agents who will suffer it if we opt to prevent harm A, and (2) that while harm C is not as serious as harm A, neither is it trivial by comparison.

Perhaps the point can be made clearer by putting it in terms of the two-level conception of practical reason. It would indeed be strange and unpalatable if the option of preventing harm A to one person served only to determine whether we ought to prevent harm C to five people, or harm E to fifty people. But this is not the case. The option of preventing harm A to one person not only provides a second-order reason not to consider preventing harm E, it also provides a first-order reason to prevent harm A. Unfortunately for the person who will suffer harm A, this first-order reason is balanced by the first-order reason to prevent harm C to one person, while the first-order reason to prevent harm C to a second person breaks the tie in favor of option 2 (preventing harm C to five people). The agent confronted by this three-option case ought to consider the possibility of rescuing the one person who will suffer harm A (or at least he ought to be

open to doing so); if he does not (or would not), then he fails to recognize the significance of the loss that one person stands to suffer. But in preventing harm C to five people, he does not fail to recognize this, or to take fully into account that agent's claim to be saved. For as Scanlon notes, the importance of saving an agent can be fully taken into account even though, because of the importance of saving other agents, no attempt to rescue the first agent ought to be made.²²

In failing to even consider the possibility of preventing harm E, regardless of how many people will suffer it, when it is possible to prevent one person from suffering harm A, does an agent thereby trivialize the loss to each of the individuals who will suffer harm E? I think not. Rather, the claim is that harm E *is* trivial in comparison to harm A, and that this is a fact that all agents, including those who will suffer harm E, ought to recognize. These agents have no complaint about how the rescuer treats them, therefore, because the rescuer simply responds to a fact that each of them can (or at least ought) to recognize. Moreover, we ought to keep in mind the limited scope of the exclusion on an agent taking into account the possibility of preventing harm E. As already noted, when confronted by a choice between harm C and harm E, an agent ought to choose on the basis of how many people will suffer the respective harms, and in some cases the possibility of preventing harm E may exclude the possibility of preventing other, lesser, harms, regardless of how many people will suffer them.

In sum, once we recognize that option 1 provides an agent with a second-order reason to exclude option 3 from his deliberation, as well as a first-order reason for action, we will see that the rescuer in Norcross's example faces two consecutive two-option cases, not a single three option one. Initially, that choice is between preventing harm A

to one person and preventing harm C to five people. That it is also physically possible for the rescuer to prevent fifty people from suffering harm E is morally irrelevant (just as the physical possibility of providing n people with the small benefit of an ice cream cone would be), and so ought not to figure in his deliberation. At some later point in time the rescuer faces a new choice between preventing harm C to five people and harm E to fifty people. Assuming that none of the agents who stand to suffer harm are responsible for the change in the situation the rescuer confronts (as would be the case if someone at risk of suffering harm E had caused the leak in the rescuer's gas tank), how a particular rescue situation comes to be seems irrelevant to the question of what the rescuer ought to do in that situation. That an agent in danger of suffering harm C would have been saved in a different rescue situation provides no basis for an objection to his not being saved in the present one.

I do not expect that the arguments of the previous few paragraphs will convince a Utilitarian such as Norcross, who will likely agree that a headache is trivial in comparison to death, but assert that the aggregate harm of very many headaches is not. Such a rejoinder returns the debate between Scanlon and Norcross, and between (most) consequentialists and non-consequentialists, to the question of whether proper moral reasoning includes interpersonal aggregation and the maximization of value. My primary aim has not been to settle this debate, however, but only to defend the coherence of the concept of a morally relevant harm as it figures in common-sense and contractualist moral theory (and perhaps other non-consequentialist moral theories as well). The orbital conception of a morally relevant harm, understood in terms of the two-level conception of practical reason sketched herein, provides such a defense.²³

¹ Any of these claims regarding the content of common-sense morality may be challenged, of course, but most of the authors participating in the debate over the concept of a morally relevant harm seem willing to concede claims of the sort made here.

² T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Belknap Press of Harvard University Press, 1998), pp. 239-240.

³ Alastair Norcross, "Contractualism and Aggregation," in *Social Theory and Practice* 28(2): 303-314. Other philosophers that have commented on Scanlon's remark concerning morally relevant harms, and the potential problems such a concept gives rise to, include F. M. Kamm, "Owing, Justifying, and Rejecting," *Mind* 111 (2002): 351-3; Derek Parfit, "Justifiability to Each Person," *Ratio* XVI:4 (December 2003): 384-88; Jussi Suikkanen, "What We Owe to Many," *Social Theory and Practice* 30:4 (2004): 499-501; and Sophia Reibetanz, "Contractualism and Aggregation," *Ethics* 108 (January 1998): 309-11.

⁴ Scanlon, *What We Owe to Each Other*, p. 240. For a more nuanced view of when, and in what ways, a less serious harm may still be morally relevant to a greater one, see F. M. Kamm *Morality, Mortality* Volume 1 (New York: Oxford University Press, 1993), p. 171.

⁵ Norcross, "'Contractualism and Aggregation,'" p. 307.

⁶ Of course, there might be more than two categories of harm, so that M and Z would be in different categories, and not relevant to one another. But the kind of problem noted in the text will arise in any case where one category contains three or more harms. So for example, even if category 2 contained only the harms M, N, and O, M and O would be

relevant to one another, while M and L would not be, despite the fact that there is less difference in harm between M and L than there is between M and O.

⁷ This sketch of a two-level conception of practical reason obviously owes a great deal to the work of Joseph Raz; see, for instance, Raz, *Practical Reason and Norms*, 2nd Edition. (Princeton: Princeton University Press, 1990).

⁸ As Norcross remarks, “contractualism departs from consequentialism in its account of precisely which reasons can balance which other reasons. For the consequentialist, any kind of morally relevant reasons can, in principle, balance any other kind. [Note: This should read “any kind of moral reason,” not “any kind of morally *relevant* reason.”] Scanlon, on the other hand, claims that some moral reasons are simply not “relevant” to other moral reasons” (Norcross, “Contractualism and Aggregation,” p. 305). The two accounts of reasons and of practical reasoning described in the text reflect this difference between consequentialism and contractualism (and many other non-consequentialist theories).

⁹ Scanlon, *What We Owe to Each Other*, p. 51.

¹⁰ For example, it might be that creatures that employ the two-level theory of practical reasoning do better at maximizing utility in the long-term than they would if they always took into account all of the harms they could prevent.

¹¹ John Rawls, *Political Liberalism* (New York, 1993), pp. 56-?. For discussion, see David Lefkowitz, ‘A Contractualist Defense of Democratic Authority’, *Ratio Juris* 18:3 (2005), pp. 349-50.

¹² Lefkowitz, ‘Contractualist Defense’.

¹³ For a defense of this claim, see David Lefkowitz, “On a Moral Right to Civil Disobedience’, *Ethics* 117 (2007), pp. 228-32.

¹⁴ At this point, the dedicated utilitarian may claim that it is at least reasonable to hold that a minor, temporary, headache falls within the orbit of death. Not surprisingly, I disagree. Nevertheless, I do not intend the point about reasonable disagreement raised in the text to rule out the utilitarian’s claim by stipulation. Rather, I seek only to add a bit more detail to the conception of a relevant reason defended in this paper, and in particular to show that it can accommodate different judgments regarding the relevance of two harms. That it can do so contributes to its plausibility. Whether a non-consequentialist moral theory that includes the orbital conception of a relevant harm ultimately proves to be superior to a consequentialist one that does not (or that treats headaches and death as relevant to one another) requires a holistic assessment of the two competing theories, which I do not attempt to offer here.

¹⁵ Similarly, there may be clear cases of being bald, and clear cases of being not bald, and cases in which there simply is no answer to the question ‘is this person bald?’

¹⁶ Like Scanlon, Norcross, and the other discussants listed in footnote 3, I focus here on rescue cases, where the probability of successful rescue is the same for all parties. That the possibility of preventing a death excludes preventing any number of headaches in this case need not entail that just institutions for allocating resources for the prevention of harm will focus solely on the prevention of death caused in various ways, regardless of the likelihood of someone dying in these ways, or the likelihood of successfully preventing it, with no attention paid to preventing headaches (or broken arms).

¹⁷ This conclusion oversimplifies Scanlon's position, in that it suggests a straightforward preference for saving the many over the few when choosing amongst relevant harms, whereas Scanlon makes only the weaker, vague, claim that in some cases of choosing amongst relevant harms, it is permissible (or perhaps even obligatory) to save the many rather than the few. Exactly what principle we ought to employ in determining which of two relevant harms we ought to prevent, when we cannot prevent both, is a crucial question, but one I cannot address here.

¹⁸ Norcross does not note this difficulty, instead simply stipulating that one of these options is obligatory. Other commentators on Scanlon's discussion of aggregation have done so, however. See, for example, Thomas W. Pogge, "What We Can Reasonably Reject," in *Philosophical Issues* 11: 118-147; Parfit, "Justifiability to Each Person," 384.

¹⁹ Norcross, "Contractualism and Aggregation," p. 309. Norcross does not discuss any particular basis, such as the orbital conception of morally relevant harm, for treating one of the three options (1, 2, or 3) as morally obligatory, but instead simply stipulates that one of three options is required. If successful, however, his objection holds regardless of the justification offered for treating one of these three options as obligatory.

²⁰ Norcross, "Contractualism and Aggregation," p. 309.

²¹ Indeed, Joseph Raz claims that "Scanlon implicitly acknowledges that Contractualism, while presupposing such a classification [of the moral importance (i.e. relevance) of different harms], contributes nothing to establish it" (Raz, "Numbers, With and Without Contractualism." *Ratio* XVI:4 (December 2003): 364-5). Scanlon disputes this claim, but also notes that "the challenge for contractualism is to overcome the objection that it rules out plausible aggregative arguments," and that the arguments he offers to meet this

challenge need not be uniquely contractualist (Scanlon, "Replies." *Ratio* XVI:4 (December 2003): 431).

²² Scanlon, *What We Owe to Each Other*, p. 234.

²³ I wish to thank Scott James for helpful discussion of the ideas in this paper, and the University of ----- for the award of a Summer Excellence Grant, which was instrumental to the completion of this paper.