

H. M. MALM

BETWEEN THE HORNS OF THE  
NEGATIVE-POSITIVE DUTY DEBATE

(Received 23 February, 1989)

Two principles dominate the literature on the moral significance of the difference between duties not to cause harm and duties to prevent harm. One maintains that duties not to cause harm are *stricter* than duties to prevent harm, while the other maintains that the duties are morally equivalent. Each principle has been defended with seemingly persuasive examples and has a prevalent intuition in its favor. Each principle has also been criticized with equally persuasive counter-examples, suggesting that some of the implications of that view are untenable. Thus we seem to be in the uncomfortable position of having to choose between the competing principles when neither one is entirely acceptable. As Michael Gorr puts it:

We are left, I think, with the view that the choice between [these principles] is a choice between *fundamental* principles. . . . [with there being] little we *can* do [to choose between them] other than compare the facility with which each accommodates the entire range of our considered moral judgments.<sup>1</sup>

I think we are not forced to choose. In this paper I develop an account of the moral significance of the difference between duties not to cause harm and duties to prevent harm that captures the strong points of each of the currently rival views while escaping their counter-intuitive implications. In the first section I clarify the terms and scope of the discussion that follows. In section II I set out our *evidence*, that is, the particular judgments and underlying intuitions that support either of the two rival views, and thus the judgments and intuitions that must be captured by a mediating view. I also argue in this section that while the difference between the duties is significant, the standard bases for explaining the difference, and the standard way of unpacking the central notion of "stricter" are incorrect. In sections III and IV I develop an account of the moral difference between the duties that

focuses on the reasons that can justify or excuse violations of the duties in relevantly similar situations. This account includes a new definition of “stricter” and a set of principles explaining the sense in which duties not to cause harm are stricter than duties to prevent harm. I conclude by defending this account as morally preferable to the others that have been discussed.<sup>2</sup>

## I

I will refer to duties not to cause harm as *negative duties*, and to duties to prevent harm as *positive duties*. Paradigmatic examples of negative and positive duties are the duty not to kill and the duty to prevent a death.<sup>3</sup> The relevant class of positive duties consists of duties to prevent harm that an agent has in “virtue of being a member of the moral community,” and not, in contrast, in virtue of a special relationship between the agent and potential victim (e.g., via a contract, promise, or one’s role as custodial parent).<sup>4</sup> These latter duties, so-called “special duties,” can exist along side positive duties and oblige the same act. (A lifeguard, for example, may have both a special duty and a positive duty to throw a rope to a drowning swimmer.) But given that duties not to cause harm do not rest on a special relationship between the agent and potential victim, duties to prevent harm that do are not part of the proper comparison class of positive duties.<sup>5</sup>

Violations of negative and positive duties are, of course, acts of causing harm and failing to prevent harm. (I will use “acts” in both cases.) However, since my focus is on the difference between the duties, I will exclude from consideration any acts of causing harm and failing to prevent harm that are not at least *prima facie* prohibited by negative and positive duties, and use “violations” in a factual, nonevaluative sense. (Thus there is nothing incoherent in the notion of a justified or permissible violation.) Also, when discussing particular violations of positive duties, I will be assuming that the agent in question had the ability, the opportunity, and the knowledge of the ability and opportunity to prevent that harm. I take these to be necessary conditions for saying that an agent ‘failed to prevent,’ or ‘refrained from preventing’ a harm. Similar assumptions will be made about violations of negative duties, in order to retain a parallel.

To ask whether there is a morally significant difference between negative and positive duties is to ask whether it *can* matter, from a moral point of view, that a given piece of behavior violates a negative duty rather than a correlate positive duty, not merely whether it does matter in one particular case. Thus the view that there is not a significant difference entails that there are no situations in which other things are equal and violations of the duties deserve different moral assessments; the acts will be both wrong, both permissible, or both right, and the agents will be equally blameworthy, praiseworthy or excused. Conversely, the view that there is a significant difference entails that there is at least one situation in which other things are equal and violations of the duties deserve different moral assessments. It is up to the particular account of the significant difference to determine whether the violations of the duties will always deserve different moral assessments (other things equal), and how those assessments differ.

The assumption that other things are equal is needed to ensure that any difference in the moral status of the acts or agents is tied to the difference between the duties, and not, for example, to a difference in risks to the agent. The “other things equal” clause will henceforth be assumed to cover (i) the evilness of the agent’s motive, (ii) the presence of a prior intention to perform the act, (iii) the amount of harm that follows from the act, (iv) the agent’s subjective certainty that the harm will ensue, (v) the effort, risk, or sacrifice required of the agent not to perform the act, (vi) the presence of any other duty that conflicts with the one in question, and (vii) the dischargeability of the sets of the agent’s negative and positive duties.<sup>6</sup> It allows for differences that are irrelevant to the assessments of the acts or agents (e.g., the color of the victim’s hair), and differences that are needed to have a violation of a negative duty on the one hand and a positive duty on the other.

The examples used in the negative-positive duty debate can be divided into two main types. The first type will be called *conflict examples*. These have one situation and one agent who cannot fulfill each of the prima facie duties incumbent upon her. For example, she may have to choose between fulfilling the duty to prevent Jones’ impending death and fulfilling the duty not to kill Smith, because the means for fulfilling the one duty entails, in the given situation, the nonfulfillment of the other. The question in these cases is whether the

agent's *alternatives* are morally equivalent, or equally, whether one of her duties overrides the other.

The second type will be called *comparison examples*. These have two situations which differ only in that the agent in one situation violates a negative duty while the agent in the other situation violates a correlate positive duty. These examples ask us to determine whether the violations of the duties deserve equivalent moral assessments in separate but similar situations. A subcategory of comparison examples is worth mentioning here. These are cases in which we are comparing two conflict situations. On one side we have a conflict between negative duty N and some other duty O, and on the other side we have a conflict between positive duty P and the same other duty O. N and P are the two duties being compared and the question is whether a violation of N, in light of conflicting duty O, deserves the same moral assessment as a violation of P, in light of conflicting duty O (e.g., are the agent's alternatives morally equivalent in one case but not in the other?).

Finally, when assessing competing principles about the moral significance of the difference between negative and positive duties I will be employing a sort of coherence theory of justification, in that the morally preferable principle will be the one that best captures our considered moral judgments about a wide range of particular cases, while being consistent with our fundamental moral beliefs and values, and the other principles within our moral network. Thus I will not be attempting to prove here that our considered moral judgments about particular cases are correct, but rather taking these judgments as data and looking for a way to explain them within a theoretical framework.<sup>7</sup>

## II

In this section I set out the two currently rival views, and the grounds for thinking that there is, and that there is not, a morally significant difference between negative and positive duties.

One side of the debate maintains that there is no morally significant difference between negative and positive duties. I will call this view the *Moral Equivalence Principle*. It entails that violations of correlate duties deserve equivalent moral assessments when other things are equal.

Proponents of the Moral Equivalence Principle are quick to admit

that actual acts of causing harm are often worse than actual failures to prevent harm. They argue, however, that this difference in the moral status of the acts can always be attributed to differences in such things as motives, consequences, and risks to the agents.<sup>8</sup> They maintain that when other things are truly equal, it will be clear that failures to prevent harm are every bit as bad as acts of causing harm.

The typical way to defend the Moral Equivalence Principle is to offer comparison examples involving only minimal effort or risk for the agent and no conflicting duty. In such cases it is assumed to be intuitively clear that the acts and agents deserve equivalent moral assessments. Consider the following pair of cases.

- (1) Smith walks into a room and discovers that a machine that has been set up to crush the child inside has malfunctioned. Smith knows that she could re-start the machine by pushing a nearby button, and she does so solely because she is curious to see how flat a person can be.
- (2) Jones walks into a room and discovers that a machine that has been set up to crush the child inside is about to do so. Jones knows that she could stop the machine by pushing a nearby button, but she does not do so solely because she is curious to see how flat a person can be.<sup>9</sup>

Smith pushed a button and killed a child, while Jones refrained from pushing a button and allowed a child to die. I will assume, as I think is correct, that our considered moral judgments are such that both acts are wrong and the agents are equally blameworthy. An agent, that is, who allows a child to die simply to satisfy her curiosity as to how flat a person can be, is deserving of as much moral disapprobation as an agent who causes a child to die for the same reason. Since these judgments accord with the dictates of the Moral Equivalence Principle, they provide evidence for that view.

In the above example Smith acted and Jones refrained. Other comparison examples support the Moral Equivalence Principle without drawing attention to (the irrelevance of) the difference between acting and refraining.

- (1) Smith is mixing some chemicals that she will use to develop photographs. She is told that if she adds the least expensive of two chemicals to the solution (each of which will do the job), a gas will

be produced that will kill a person in the next room. Chemical A costs \$1.00 and chemical B costs \$2.00. Smith decides that she'd rather not spend the extra dollar, and for this reason she adds chemical A to the solution and the person in the next room dies.

- (2) Jones is mixing some chemicals that she will use to develop photographs. She is told that if she adds the more expensive of two chemicals to the solution (each of which will do the job), a gas will be produced that will save the life of a person in the next room. Chemical A costs \$1.00 and chemical B costs \$2.00. Jones decides that she'd rather not spend the extra dollar, and for this reason she adds chemical A to the solution and the person in the next room dies.

Though Smith and Jones moved in the same way, Smith caused a death and Jones failed to prevent a death. Each did so solely to avoid spending an extra \$1.00.<sup>10</sup> I will assume, again as I think is correct, that our considered moral judgments are such that both acts are wrong and the agents are equally blameworthy.

Similar examples abound. The Moral Equivalence Principle captures the prevalent intuition that an agent *can be* as morally blameworthy for failing to prevent a death as she would have been had she killed. It also accords with the more specific claims (supported by the above examples) that causing a harm is not *in itself* worse than failing to prevent a harm, and that the distinction between acting and refraining is not in itself a morally important matter. (I will discuss the difference between these claims later.) Its strongest support comes from our judgments in comparison examples involving only minimal effort or risk for the agents and no conflicting duty.

The main problem for the Moral Equivalence Principle is that while it accords with our judgments in at least one sort of case, it fails with respect to others. These are the cases that draw us towards the opposing view.

The opposing view is captured in what I will call the *Intrinsic Difference Principle*. It asserts that duties not to cause harm are *stricter* than duties to prevent harm, where "stricter" is defined as follows:

Duty S is *stricter* than duty T if and only if, other things equal, an agent is morally more culpable for violating S than [s]he is for violating T.<sup>11</sup>

The Intrinsic Difference Principle entails that it is worse to violate a

negative duty than it is to violate a correlate positive duty when other things are equal. There are at least two sorts of cases that support this principle directly, and two more sorts of cases that support the general claim that there is *some* morally significant difference between negative and positive duties.

First, the Intrinsic Difference Principle derives support from our judgments in conflict examples. Suppose that Jones is in dire need of a heart transplant and that the doctor has the following alternatives: She can kill another patient, Smith, transplant Smith's heart to Jones and thereby prevent Jones' death, or she can not kill Smith and thereby let Jones die.<sup>12</sup> Our considered moral judgments indicate, I assume, that the doctor's alternatives are not morally equivalent — she ought not to kill Smith even though this requires her to allow Jones to die.<sup>13</sup> It may be objected that the one alternative is preferable only because the other alternative involves using Smith as a means. But that is not a necessary condition for a moral difference in a conflict situation. Suppose that a treatment needed to prevent Jones' death would produce, as a by product, a gas that would kill Smith in the next room. It still seems clear that the doctor's alternatives are not morally equivalent.<sup>14</sup>

The Intrinsic Difference Principle accords with this judgment because the claim that it is worse to violate a negative duty than it is to violate a correlate positive duty entails that in a conflict situation, other things equal, the duty not to kill overrides the duty to prevent a death.

Second, the Intrinsic Difference Principle derives support from our judgments about comparison examples involving a significant harm or loss for the agent. Suppose that in the photographic-solution example chemical A costs \$1 and chemical B costs \$1,000. Or consider the following pair of cases involving a significant physical harm.

- (1) Smith is a professional quarterback. While driving down an icy hill he notices a man passed out in the middle of the road. Smith knows that he must quickly apply his brakes if he is to avoid running over and killing this man. He also knows that if he applies his brakes on the icy hill his car will swerve into a ditch and the crash will cause permanent injury to his legs and force him to use a cane for the rest of his life. Smith does not want to injure his legs, and for this reason he runs over and kills the man.
- (2) Jones is a professional quarterback. While driving down an icy hill he notices a man bleeding to death on the roadside. Jones knows

that applying his breaks on the icy hill would cause his car to swerve into a ditch and that the crash would cause permanent injury to his legs and force him to use a cane for the rest of his life. He also knows that he would still be able to save the man's life. Jones does not want to injure his legs, and for this reason he drives on and the man dies.

The factual differences between these cases are ones that allow for the presence of a negative duty on the one hand and a positive duty on the other. The important similarity is that the agent in each case chose to protect his own welfare at the cost of another person's life. With respect to the moral status of their acts, it is clear that Smith's act was wrong and he is morally blameworthy: The risk of harm to oneself does not justify the infliction of an even greater harm on an innocent other. With respect to Jones' act, there may be controversy about its permissibility. Some might argue that the harm to Jones, were he to help the man, was sufficiently great to permit his failure to do so. Others might argue that it was not sufficient in this case but would have been sufficient had the harm been, say, permanent paralysis. We need not resolve this controversy here. Insofar as we would judge Jones to be at least less blameworthy than Smith (and I think we would), we recognize some morally significant difference between the negative duty not to kill and the positive duty to prevent a death. The Intrinsic Difference Principle accords with this judgment by implying that an agent can be required to go to greater lengths in order to avoid causing a harm than would be required of her to prevent a similar harm.

The third sort of case involves a comparison of two conflict situations. On one side we have a conflict between negative duty N and a correlate duty O, and on the other side we have a conflict between positive duty P and O.

First recall the previous conflict examples in which the means for preventing Jones' death would kill Smith (i.e., the transplant example and gas-production examples). In each case the agent faced a conflict between a negative duty not to kill (here duty N) and a positive duty to prevent a death (here duty O). I assumed that our considered moral judgments are such that the agent's alternatives are *not* morally equivalent. Now consider a case in which the agent faces a conflict between a positive duty to prevent a death (duty P) and another positive duty to

prevent a death (duty O). For example, suppose that a doctor has two patients, each of whom need a heart transplant, and only one heart is available. In this case, unlike the previous cases, the agent's relevant alternatives *are* morally equivalent: Provided that she saves one of the two patients, it is morally indifferent which one she saves. This change in our considered moral judgments from one conflict to the next supports the recognition of a morally significant difference between N and P.

The final sort of case also involves a comparison of two conflict situations but now the conflicting duties are non-similar, that is, they involve different levels of harm. Suppose that an agent has to choose between causing a harm of degree X to Smith, and causing a harm of degree  $X + n$  to Jones. Or suppose that an agent has to choose between failing to prevent a harm of degree X to Smith and failing to prevent a harm of degree  $X + n$  to Jones. In cases in which the conflicting duties are both negative or both positive, the "size of the harm" provides an acceptable guide to resolving the conflict: other things equal, the agent ought to choose the alternative that results in the lesser harm. But when the conflict is between a negative duty and a positive duty, the "size of the harm" does not provide an acceptable guide. For while there are cases in which an agent is justified in causing a harm to one person in the efforts to prevent a greater harm to another (e.g., running over Smith's little toe in order to rush heart-attack victim Jones to the hospital), the justification for this does not rest on the mere fact that the one harm was greater. Of central importance is the size of the *disparity* between the harms. It is not permissible, I should think, to inflict serious brain-damage on Smith in the efforts to prevent only slightly greater brain-damage to Jones, nor permissible to crush one of Smith's toes just to prevent Jones from suffering two crushed toes. Also relevant is the severity of X itself and, perhaps, the effects of the particular harm on the particular recipient (though these subjective considerations may be relevant in negative-negative and positive-positive conflicts as well). It is at least arguable that one may not inflict moderate brain-damage on Smith in the efforts to prevent severe brain-damage to Jones (or even Jones' death), and arguable that one may not crush Mikhail Baryshnikov's toes in the efforts to minimize the damage of Jones' heart-attack, which we know to be nonfatal.

In general, it seems that an agent is justified in causing a harm of degree  $X$  to  $A$ , in the efforts to prevent a harm of degree  $X + n$  to  $B$  only if either (a) the disparity between  $X$  and  $X + n$  is great, and  $X$  is not itself quite serious, or (b)  $A$  can and will be compensated for the harm she suffers, at least to a degree that makes that harm consistent with (a). If something at least similar to this is the case, then while the “size of the harm” is an acceptable guide for resolving nonsimilar negative-negative conflicts, and nonsimilar positive-positive conflicts, it is not an acceptable guide for resolving nonsimilar negative-positive conflicts.

In summary, I have discussed four sorts of cases in which our considered moral judgments support the view that there is some morally significant difference between negative and positive duties, and one sort of case in which they support the view that there is not. Corresponding to these cases are the following five general claims (underlying intuitions): (i) the duty not to kill overrides the duty to prevent a death, (ii) an agent can be required to go to greater lengths in order to avoid causing a harm than would be required of her to prevent a similar harm, (iii) an agent’s alternatives are morally equivalent in similar negative-negative conflicts and similar positive-positive conflicts, but not in similar negative-positive conflicts, (iv) the “size of the harm” provides an acceptable guide for resolving nonsimilar negative-negative conflicts, and nonsimilar positive-positive conflicts, but not for resolving nonsimilar negative-positive conflicts, and (v) an agent can be as morally blameworthy for failing to prevent a harm as she would have been had she caused that harm.

In light of these cases and underlying intuitions I contend that the relevant issue is not whether there *is* a morally significant difference between negative and positive duties, for the evidence in favor of a difference is, I believe, overwhelming. Rather, the relevant issue is determining just what that difference is. In other words, given that the difference between duties not to cause harm and duties to prevent harm is morally significant, and given that the former are in some sense stricter than the latter, how are we to unpack the central notion of “stricter” within a moral principle that accounts for our judgments and underlying intuitions?

The Intrinsic Difference Principle has been presented as the stan-

standard answer to this question. It directly accords with claims (i) and (ii), and indirectly accords with claims (iii) and (iv) (it is indirect because the principle itself is silent with respect to conflicting duties of the same type). Its main problem, for our present concerns, is that it is inconsistent with claim (v) — a claim well supported by the kinds of comparison examples offered on behalf of the Moral Equivalence Principle. Before turning to an alternative account I will briefly sketch what seems to me to be the source of the problem.<sup>15</sup>

The minimally sufficient condition for a morally significant difference between negative and positive duties is one case in which other things are equal and violations of the duties deserve different moral assessments. We have already seen that there are such cases. In various defenses of the Intrinsic Difference Principle we attempt to account for the difference in the moral status of the acts by attributing it to some “in itself” significant difference in the natures of the acts. In other words, we seem to have first, an assumption that if a particular act of causing harm is worse than a particular failure to prevent harm when other things are equal, then there must be some difference in the properties of the act-types causing harm and failing to prevent harm that makes the one particular act worse than the other, and second, arguments about what that difference is. Some accounts focus on the difference between acting and refraining (arguing that *acting* to cause a harm is in itself worse than intentionally *not acting* to prevent a harm), other accounts focus on the rights of the victim (arguing that acts of causing harm violate rights, failures to prevent harm do not, and that doing the former is in itself worse than doing the latter), and still others on the responsibility of the agent.<sup>16</sup> While each account has problems of its own (some of which I discuss later), they seem to share the problem of implying that particular acts of causing harm will be worse than particular failures to prevent harm whenever other things are equal. For if, as it is being claimed about our evidentiary cases, the one particular act is worse than the other *simply because* the one has property A and the other has property B, and if A and B are properties of the act-types and thus present in every instance of the acts, then the difference between A and B will effect a difference in the moral status of the acts in every pair of cases that differ only in that way.<sup>17</sup> Thus insofar as we attribute the difference in the moral status of the particular acts to some

in itself significant difference between the properties of the act-types — one that makes acts of causing harm *in themselves worse* than failures to prevent harm — we will find ourselves unable to account for the set of intuitions (i)—(v).

### III

What we need is an account of the moral difference between the duties that does not entail that violations of negative duties are always worse, other things equal (nor in themselves worse), than violations of correlate positive duties, and thus is able to account for our judgments and underlying intuitions in all the discussed cases. The basis for this account has already been implicitly established. A review of claims (i)—(iv), the claims entailing that negative and positive duties are morally *different*, shows that they all relate to cases in which the negative or positive duty in question *conflicts* with some other morally relevant consideration — be it the opposing negative or positive duty as in standard conflict examples, a risk of harm to the agent, or a third duty that conflicts with each of the other two. Claim (v) on the other hand, the claim entailing that violations of negative duties are not always worse, other things equal (nor in themselves worse), than violations of correlate positive duties, is supported by cases in which there is no morally relevant consideration that conflicts with the agent's duty and thus nothing that can count in favor of a violation of that duty. What this suggests is that the significant difference between the duties is a difference in the *reasons* that the correlate duties admit as grounds for justified or excused violations. In other words, positive duties permit or excuse violations on grounds and (or) in circumstances that correlate negative duties do not. This sort of difference is captured in the following preliminary principle:

DR — Correlate negative and positive duties are morally *different* if it is possible that there be two situations which differ only in that situation 1 contains negative duty N and situation 2 contains correlate positive duty P, and there is present in both situations a reason R whose justificatory or excusatory effect with respect to a violation of N is different from its justificatory or excusatory effect with respect to a violation of P.

DR provides a set of evidentiary conditions for a morally significant difference between general negative and positive duties. By “general

duties” I mean the general rules or moral prohibitions (e.g., ‘Do not kill’) that give rise to particular duties in particular situations (e.g., ‘do not kill Smith’). If fulfilled, DR allows us to attribute a difference in the moral status of particular acts of causing harm and failing to prevent harm to a difference in the general duties proscribing those acts: The acts deserve different moral assessments, when they do, because the general duties are such that one duty permits or excuses violations on grounds and (or) in circumstances that the other does not.<sup>18</sup>

In what follows I discuss two overlapping directions that one might take to develop and explain this sense of a morally significant difference. These accounts are based, respectively, on a difference in the strength of reasons, and a difference in the types of reasons that can justify or excuse violations of correlate duties when other things are equal. For brevity, I will concentrate on the justificatory effect of reasons.

The first way to develop DR is based on the idea that two general duties are significantly different if one of them can require stronger reasons in order to be justifiably violated than would be required by the other duty in a similar situation. To develop this view I need to give brief accounts of the notions of a “relevant” reason and a “stronger” reason, as these apply to reasons for performing an otherwise prohibited act.

**Relevance:** A reason of the form ‘doing X would involve or result in O’ is *relevant* to assessing whether an agent is justified in violating a general duty that requires act X in circumstances C, if and only if (i) the reason has some moral weight in support of a violation of that duty and (ii) it applies to C.

For example, suppose that the duty in question is the duty not to lie, such that ‘doing X’ is telling the truth. With respect to condition (i), the reason “telling the truth would involve breaking a promise” is a reason with some moral weight in support of a violation of the duty not to lie, while the reason “telling the truth would involve interacting with someone with blond hair” is not. With respect to condition (ii), although the reason “telling the truth would involve breaking a promise” is a reason with moral weight, it does not apply to C if telling the truth would not involve breaking a promise in that C.

The strength of a relevant reason is its *degree* of moral weight or

justificatory potential. For our present purposes we need not suppose that we have a list which ranks the exact strength of every relevant reason for every violation of a given duty. We need only recognize that some relevant reasons carry more moral weight than others in otherwise identical situations. For example, in a situation involving the duty not to lie, considerations about causing severe harm to another yield reasons with a greater strength than considerations about hurting someone's feelings. With this in mind, "stronger" can be defined as follows.

Reason R is *stronger* than reason S if and only if there are two situations, C1 and C2, which both include a duty that requires act X and are otherwise similar except that C1 involves R but not S, and C2 involves S but not R, and R is sufficient to justify the nonperformance of X in C1, while S is not sufficient to justify the nonperformance of X in C2, even though both R and S are relevant.

The important points are that the strength of a relevant reason determines whether a violation of a given duty is justified, and that some relevant reasons are stronger than others. With a certain change in focus, these points can now be used to develop DR.

The change in focus is that instead of comparing the effects of two reasons on the assessments of violations of one duty, (as in the above definition of "stronger"), we need to be comparing the effects of one reason on the assessments of violations of two duties. In other words, we need to be comparing two situations which differ only in that situation 1 contains negative duty N, and situation 2 contains correlate positive duty P, and both situations contain relevant reason R. Now if it is possible that there is an R that is sufficient to permit a violation of P in situation 2, but not a violation of N in situation 1, then N and P are not morally equivalent because in some relevantly similar situations violations of those duties are not equally permissible. Stated as a general principle we have:

DS — Two general duties are significantly different if a difference in the moral status of violations of those duties can be effected by the strength of a reason that is relevant to both violations.

The previously discussed examples and judgments indicate that negative and positive duties do differ in this way. For example, we have seen that in cases involving a significant risk to the agent, this risk can permit a failure to prevent an even greater harm to an innocent other,

in situations in which it would not permit the causing of an even greater harm. Similar claims can be made about comparisons of two similar conflict situations and two nonsimilar conflict situations. In at least some instances of each type, the conflicting duty O provides a reason that unequally affects the moral status of the acts.<sup>19</sup>

Yet unlike the Intrinsic Difference Principle, the difference described in DS is consistent with holding that violations of correlate duties sometimes deserve *equivalent* moral assessments, other things equal. In comparison examples involving only minimal effort and risk for the agent and no conflicting duty, there is no morally relevant reason in support of either violation, and thus nothing that can justify those violations or mitigate the blameworthiness of the agents. Here DS allows that the acts are both wrong and the agents are equally blameworthy.<sup>20</sup>

The second way to develop DR is to show that there is a type or class of reasons whose justificatory (or excusing) potential changes with respect to acts of causing harm and failing to prevent harm. By “change in potential” I mean, for example, that while reasons of this type are not even *relevant* when assessing violations of the duty not to kill, they are relevant when assessing violations of the duty to prevent a death. They are potential justifications that may, in a given case, be sufficient to permit a violation of that duty.

The type of reason with which I am here concerned can be most easily characterized by first mentioning two other types of reasons that *are* standardly accepted as potential justifications for killing. First, there are reasons that can be characterized by their primary reference to a fact about the victim alone. Reasons of consent or request fall in this class, as in the case of the wounded soldier who asks to be shot in the fields rather than risk capture by the enemy. Second, there are reasons that can be characterized by their primary reference to a fact about the *relationship* between the victim and another. Reasons of self-defense and defense of third parties fall in this class. In contrast, I am here concerned with a type of reason that can be characterized by its primary reference to a fact about the agent alone. Reasons based on an agent’s personal or religious convictions fall in this class, as when an agent explains that she violated a given duty because the act required to fulfill that duty was prohibited by her fundamental religious beliefs.

Generally speaking, our moral and legal practices indicate, on the

one hand, that an agent is justified in killing another only if there is something about the victim (either alone or in relation to another) that establishes that an otherwise objectionable interference with her basic rights or liberty is permissible. On the other hand, failures to prevent a death can, it seems, sometimes be justified by a fact about the agent which establishes that in *her particular case*, the standard duty to aid may be permissibly infringed.

To see this, consider the following example in which an agent fails to prevent a death in the efforts to uphold her fundamental religious convictions.

Smith is a Christian Scientist who has been told that she is the only suitable donor for a bone marrow transplant which a stranger urgently needs to live. Smith does not want the stranger to die, but her religious convictions prohibit her from consenting to any kind of surgery, even to save her own life. Smith refrains from the donation and the stranger dies. As her reason for failing to prevent the death, Smith cites the fact that her act (the non-donation) was required by her fundamental religious convictions.

Now contrast this case with a few cases in which an agent causes a death in the efforts to uphold her fundamental convictions. Consider an agent who ceremoniously kills another, not because she wants that particular person dead, but because her religious convictions require human sacrifice. Or consider a political terrorist who kills a hostage, or a person who bombs an abortion clinic knowing that an innocent person would be killed. Again, it is not so much that the agents want those people dead, but rather that their deaths are known consequences, given the situations, of the agents' upholding their convictions.

In each of these four cases an avoidable death occurred as a result of the agent acting in accordance with her fundamental convictions. The specific content of each conviction is not relevant here.<sup>21</sup> The relevant issue is whether a reason based on an agent's desire to uphold her convictions (or her interest in upholding her convictions) — whatever those convictions may be (in other words, a reason of the form "doing X would require me to act contrary to my convictions") — has the same potential for justifying a violation of a negative duty as it has for justifying a violation of a correlate positive duty. I think that our judgments about the above four cases indicate that it does not.

In cases such as the last three, it should be clear that regardless of the depth of an agent's convictions, reasons based on the desire to

uphold them are never sufficient to justify killing another. It may be that these reasons simply have no moral weight with respect to acts of causing harm, such that they are not even *relevant* when assessing violations of negative duties, or that while they have some moral weight (e.g., they may be able to justify causing a minor harm), they have so little justificatory potential that they are never sufficient to justify killing another.

Yet in cases such as the first one, cases in which an agent fails to prevent a death in the efforts to uphold her fundamental convictions, reasons based on these convictions function as at least potential justifications. The cost to the agent, were she required to act contrary to her convictions, provides a morally relevant reason that may, in a given case, be sufficient to permit a violation of her duty. Given this, the justificatory potential of these reasons can change when they are offered as reasons for violating a negative duty on the one hand and a positive duty on the other.

It may be objected that while reasons of this type are never sufficient to justify killing another, neither are they ever sufficient to justify a failure to prevent a death. The reasons function, the objection continues, at most as excuses; they may relieve the agent of blame, say, for not donating needed blood, or needed bone marrow, but they would not show that the agent's act was permissible.

This objection has some force. In response it could be argued that due to the depth to which some persons hold their religious, political or other convictions, a transgression of those convictions would be as serious a harm to the agent as would be the loss of a limb. It may leave the agent emotionally devastated, with a complete loss of self-respect, or with a debilitating sense of doom. And since the risk of serious harm to the agent is standardly regarded as a justification, and not merely an excuse, for not rendering aid, then reasons based on these convictions are potential justifications for failing to prevent a death. Further, this is consistent with holding that these reasons do not have the same potential with respect to acts of killing, because the risk of serious harm to the agent does not justify the infliction of an even greater harm on an innocent other.

Though I am more inclined to regard these reasons as potential justifications than as merely potential excuses, we need not settle the

matter here. Since DR covers the excusatory effect of reasons as well as the justificatory effect, we have evidence that DR is fulfilled if reasons based on an agent's desire to uphold her convictions can excuse to a greater degree an agent who fails to prevent a death than an agent who kills. And this does seem to be the case. Smith, if not justified, was certainly less blameworthy for her failure to prevent a death than were the other agents for their killings.

In summary, reasons based on an agent's desire to uphold her convictions are examples of a general type of reason that is characterized by its primary reference to a fact about the agent alone. Due to this characteristic (and assuming our general moral and legal practices are correct) the justificatory and (or) excusatory potential of these reasons can change when they are offered as reasons for violating a negative duty on the one hand and a positive duty on the other. Stated as a general principle we have:

DT — Two general duties are significantly different if there is a difference in the types of reasons that have a potential to justify or excuse violations of those duties, or if the degree of justificatory or excusatory potential of reasons of a certain type changes with respect to violations of those duties.

#### IV

The overlap between DS and DT need not concern us here. They are auxiliary principles helping to explain the sense in which negative and positive duties can differ with respect to the reasons they admit as grounds for justified or excused violations. When combined with DR, and the preceding discussions, they suggest the following definition of "stricter":

Duty N is *stricter* than Duty P if and only if (a) there are some reasons that can justify (or excuse) a violation of P that cannot also justify (or excuse to the same degree) a violation of N, and (b) there are no reasons that can justify (or excuse) a violation of N that cannot also justify (or excuse to the same degree) a violation of P, when (c) the violations occur in relevantly similar situations and other things are equal.

We have already seen that correlate negative and positive duties differ in accordance with clause (a). And while I have not argued that they also differ in accordance with clause (b), this does seem to be the

case. The cases in which an act of killing, for example, would be permissible while an act of letting die would not, or the cases in which an agent who kills would be less blameworthy than an agent who lets die, are cases in which other things are not equal. For example, the killing would result in a quick and painless death, the letting die a slow and tortuous one, or the victim consents to being killed but not to being allowed to die.<sup>22</sup>

Given the above definition of “stricter” let us call the claim that negative duties are stricter than positive duties the *Stricter Prohibition Principle*. As does the Intrinsic Difference Principle, it entails that violations of negative duties sometimes deserve more severe moral assessments than violations of correlate positive duties, other things equal. But unlike the Intrinsic Difference Principle, it attributes the difference in the moral status of the acts to a difference in the general duties proscribing those acts, and not to some underlying difference that makes acts of causing harm in themselves worse than failures to prevent harm. This is important for at least three reasons.

First, it allows the Stricter Prohibition Principle to avoid suggesting that the difference between acting and refraining is in itself a morally important matter. This is important, not only because of the cases indicating that acting is at least not in itself *worse* than refraining, but because the difference between acting and refraining does not even correlate with the difference between negative and positive duties. Consider the following example.

Sue is in her boat on the river. Due to the current she can only row downstream or to the right. She is done fishing for the day and is about to row to shore, when she looks upstream and notices a boat capsized in the distance and a tired person, Mary, swimming towards her. Sue realizes that if she refrains from rowing for ten more minutes, Mary will be able to swim up to the boat and grab onto the side. This will prevent Mary's death. (Sue does not need to row Mary into shore since a short rest will enable Mary to swim in on her own.) Sue also realizes that if she rows into shore now, or even downstream, Mary will surely drown. Sue rows to shore and Mary dies.

In this case the positive duty to prevent a death obliged the agent to refrain from a positive act rather than (as is more commonly the case) perform a positive act, and was violated by acting rather than refraining. Now consider the following example in which the negative duty not to kill obliges the agent to act rather than refrain, and is violated by refraining rather than acting.

Sue is in her boat on the river. Quite some time ago she decided to quit rowing for a while so that she could read and sunbathe. Now, as she looks up from her book, she notices that her boat is drifting towards John who is floating peacefully on a sailboard in front of her. Sue realizes that unless she rows to the right or to the left, her boat will crash into John and he will die. Sue does nothing, John is struck and he dies.

More could be said about the difference between acting and refraining. The point here is that negative and positive duties can each oblige an agent to act as well as refrain, and can each be violated by refraining as well as acting. Thus any account of the moral difference between the duties that puts weight on the difference between acting and refraining is destined for trouble. The Stricter Prohibition Principle avoids this trouble.

Second, the Stricter Prohibition Principle accommodates the main insight, but avoids one of the problems of a rights-based defense of the Intrinsic Difference Principle. According to this defense, acts of causing harm violate rights, failures to prevent harm do not, and it is this difference, rather than the difference between acting and refraining, that makes acts of the one sort morally worse than acts of the other sort.<sup>23</sup> As it stands, this defense is problematic on two grounds. First, it is at least arguable that failures to prevent harm do violate rights — rights that do not arise from a special relationship between the agent and potential victim.<sup>24</sup> I will not pursue this further. Second, even if failures to prevent harm do not violate rights and acts of causing harm do, the significance of this difference is not one that makes acts of the one sort morally worse than acts of the other sort. We have already seen that an agent can be as blameworthy for an act that does not violate a right as for an act that does (re: the machine example), and others have offered arguments to this end.<sup>25</sup>

The structure of the Stricter Prohibition Principle allows one to maintain that the difference between violating a right and not violating a right is morally significant, without cashing out this significance in terms of “morally worse.” The presence of a right on the one hand and not on the other may be part of the reason why some conflicting considerations provide reasons which unequally affect the moral status of violations of negative and positive duties. When the conflicting consideration derives from a right of the agent, we would have, on the one hand, a conflict between two rights (one of the victim and one of the

agent) and on the other hand, a conflict between, say, an interest and a right. Or when the conflicting consideration derives from some societal good or utilitarian aim, we would have, on the one hand, a conflict between a right and this aim, and on the other hand, an interest and this aim. But holding that these conflicts may be resolved in different ways because the presence of a right prevents certain sorts of considerations from providing the justification that they would have provided were a right not present,<sup>26</sup> does not commit us to holding that in the absence of such considerations an act which violates a right is morally worse, in and of itself, than an act which does not. In short, the Stricter Prohibition Principle denies a necessary connection between “morally significant” with respect to a conceptual difference, and “morally worse” with respect to acts that bear that difference.

Finally, by defining “stricter” in terms of the reasons that can justify or excuse violations of correlate duties, the Stricter Prohibition Principle accords with judgments and underlying intuition that support the Moral Equivalence Principle. For holding that violations of positive duties can sometimes be justified or excused by reasons that would not have the same effect on violations of correlate negative duties (other things equal), is consistent with holding that in circumstances in which there is no morally relevant reason in support of either violation, the acts are both wrong and the agents are equally blameworthy.

Thus the Stricter Prohibition Principle not only captures the judgments and underlying intuitions that support the Intrinsic Difference Principle, but also those that support the Moral Equivalence Principle. And in so doing, it provides a path between the horns of the negative-positive duty debate.<sup>27</sup>

#### NOTES

<sup>1</sup> Michael Gorr, “Some Reflections on the Difference Between Negative and Positive Duties” in *Tulane Studies in Philosophy: Positive and Negative Duties*, ed. Eric Mack, 1985. Though Gorr seems to assume that these principles are the only tenable options, I am assuming only that they are the two most popular options. Elements of a third view are suggested in some of Philippa Foot’s articles. See note 2.

<sup>2</sup> The account I develop shares some points with Philippa Foot’s views in “The Problem of Abortion and the Doctrine of the Double Effect” (*Oxford Review*, no. 5, 1967, reprinted in *Killing and Letting*, ed. Bonnie Steinbock, Prentice Hall, Inc, Englewood Cliffs, N. J., 1980), and “Euthanasia,” (*Philosophy & Public Affairs*, vol. 6, no. 2, 1977). In particular, Foot maintains, as do I, that duties not to cause harm are

*stricter* than duties to prevent harm, while denying that acts of causing harm are in themselves worse than failures to prevent harm, and she focuses on a difference in the permissibility of the acts as evidence of a moral difference between them. But Foot does not provide a definition of “*stricter*” that enables us to make sense of these points. One of the purposes of this paper is to provide such a definition. For a discussion of some of the problems in Foot’s own account see Ann Davis, “The Priority of Avoiding Harm” in *Killing and Letting Die*.

<sup>3</sup> In using the terms “negative” and “positive” I am not assuming any necessary correlation between negative and positive duties on the one hand, and negative and positive acts on the other. Negative duties oblige persons to abstain from harming others, and positive duties oblige us to provide aid. I take it to be logically possible, however, that means for fulfilling the duties can include both negative and positive acts and that each type of duty can be violated by the performance of negative and positive acts. I return to this point later.

<sup>4</sup> Eric Mack, Introduction to *Tulane Studies in Philosophy: Positive and Negative Duties*, p. 1. Mack makes this distinction in terms of *general* duties, which “bind all members of the moral community to all others simply in virtue of whatever it takes to make them members of the moral community”, and *special* duties, which “exist for particular individuals in virtue of the specific and voluntary relationships into which they have respectively entered.” Though I employ Mack’s distinction, I avoid his terminology because I will later use the term “general duties” in a different way.

<sup>5</sup> Some include special duties in the relevant class of positive duties. Foot, for example, offers the “duty to look after one’s children or aged parents” as a paradigmatic example of a positive duty (“The Problem of Abortion” p. 162). Others use special duties as a tool for showing a significant difference between negative and (what I am calling) positive duties. See Raziel Abelson, “To Do or Let Happen”, *American Philosophical Quarterly*, vol. 19, no. 3, 1982.

<sup>6</sup> Raziel Abelson discusses the first five conditions in “To Do or Let Happen”. The sixth is needed to cover a comparison of two conflict situations (which I discuss in a moment), and the seventh is discussed in a number of places, including Richard Trammell, “Saving Life and Taking Life,” *Journal of Philosophy*, LXXII, no. 5, 1975 pp. 131–137, and Richard Brook, “Dischargeability, Optionality, and the Duty to Save Lives,” *Philosophy & Public Affairs*, vol. 8, no. 2, 1979. For a different view about the “other things equal” clause see Francis Myrna Kamm, “Killing and Letting Die: Methodological and Substantive Issues,” *Pacific Philosophical Quarterly*, 64, 1983.

<sup>7</sup> The judgments I cite are ones the literature shows to be widely held. When there is disagreement about a particular case it tends not to be disagreement about the proper assessment for that case, but about the implications of that assessment for a given principle.

<sup>8</sup> See, for example, Michael Tooley, “An Irrelevant Consideration: Killing Versus Letting Die,” in *Killing and Letting Die*, Bruce Russell, “On the Relative Stringency of Negative and Positive Duties,” *American Philosophical Quarterly*, 14, no. 2, 1977 (also in *Killing and Letting Die*), and Raymond Belliotti, “Negative and Positive Duties,” *Theoria*, no. 47, 1981.

<sup>9</sup> This is a comparison version of a conflict example offered by Michael Tooley (“Killing Versus Letting Die” p. 60). Judith Lichtenburg offers a similar comparison example in “The Moral Equivalence of Actions and Omissions,” *Canadian Journal of Philosophy*, Supplementary Volume VIII, 1982, p. 25.

<sup>10</sup> Notice that none of the agents, in this or the previous example, had a prior intention to bring about a person’s death, nor an especially heinous motive (which is not to say their motives were beyond reproach). This should allow the examples to escape Trammell’s charge of having a “sledgehammer” effect on our intuitions. (Trammell, “Saving Life and Taking Life”, reprinted in *Killing and Letting Die*, p. 167.)

<sup>11</sup> Bruce Russell, "On the Relative Stringency of Negative and Positive Duties" p. 217. Though Russell denies that negative duties are stricter than positive duties, his definition is both plausible and in accordance with the widespread assumption that if one duty is stricter than another, then a violation of it is worse than a violation of the other. Other definitions that may seem initially attractive are problematic. For example, if we define "stricter" in terms of overridingness, then the resulting principle is inapplicable to comparison examples (which, as we will see, can be cases in which the duties seem different). And if we weaken Russell's definition to say that an agent is only *sometimes* more blameworthy, then the resulting principle is practically useless. It fails to tell us when, much less why, one agent is more blameworthy than the other.

<sup>12</sup> This is a shortened version of an example Daniel Dinello offers in "On Killing and Letting Die" (*Analysis*, vol. 31, 1971, reprinted in *Killing and Letting Die*, p. 130). No weight is being placed on a doctor's special duty to her patients. As it stands, the doctor has this duty to both Smith and Jones, and we could just as easily imagine that the agent is a medically knowledgeable bystander but not a doctor.

<sup>13</sup> Rachels objects that the doctor would not allow Jones to die because no heart was available to save him — Smith was using his. But if the mere use of an item, or even the use of an item by the person to whom it belongs, renders that item unavailable, then we should have to say, quite implausibly, that the doctor would not allow Jones to die in the following case. The doctor needs a knife to perform a tracheotomy that will save Jones' life. Smith is using the only knife in the vicinity to cut a steak and it is Smith's knife. The doctor does not take the knife from Smith and Jones die for the lack of a tracheotomy. It seems that the initial plausibility of Rachels' objection was due to the fact that taking the heart would kill Smith — something the doctor ought not do even to save Jones. (James Rachels, "Killing and Starving to Death," *Philosophy*, 54, 1979, p. 170.)

<sup>14</sup> This is a version of one of Foot's examples in which we could produce a gas that would save five lives but create fumes lethal to one (Foot, "The Problem of Abortion," p. 163).

<sup>15</sup> I discuss the following view in more detail in "Directions of Justification in the Negative-Positive Duty Debate," *American Philosophical Quarterly*, vol. 27, no. 3, 1990.

<sup>16</sup> See, for example, Raziell Abelson, "To Do or Let Happen," Michael Wreen, "Breathing a Little Life into a Distinction," *Philosophical Studies*, vol. 46, 1984, and O. H. Green, "Killing and Letting Die," *American Philosophical Quarterly*, vol. 17, no. 3, 1980.

<sup>17</sup> There is an exception to this claim. Suppose that one agent caused a harm, another agent failed to prevent a harm, other things are equal, and both agents are insane. Here the insanity of the agents would cancel (or preclude) any difference in blameworthiness that would have been effected by the difference between causing harm and failing to prevent harm were (a) causing harm in itself worse than failing to prevent harm, and (b) both agents not insane. Cases involving cancelling considerations are not the sort of cases that are offered in defense of claim (v), and thus not the sort of cases that cause real problems for the Intrinsic Difference Principle. I will not discuss them further.

<sup>18</sup> This account includes a change in perspective from the typical defenses of the Intrinsic Difference Principle. The difference in the moral status of the acts is a consequence of the difference between the duties, and not of some "in itself" significant difference in the properties of the act-types. (The difference between the act-types is significant, if it is, *because* the difference between the duties is significant.) Further, though I develop it elsewhere, the difference between the duties is to be explained in terms of the fundamental values of the moral system in which the duties operate. See my "Directions of Justification in the Negative-Positive Duty Debate."

<sup>19</sup> Conflict examples can be accounted for in a similar way, but the explanation is more

complicated. See my "Killing, Letting Die, and Simple Conflicts," *Philosophy & Public Affairs*, vol. 18, no. 3, July 1989.

<sup>20</sup> In other words, the present view allows that failures to prevent harm are *prima facie* as bad as acts of causing harm, but denies that being 'prima facie as bad as' is equivalent to being morally indistinguishable. The acts are distinguishable in that they are violations of morally different duties, but the difference between the duties (in terms of justifying and excusing reasons) does not *require* a difference in moral status.

<sup>21</sup> To claim that the content is relevant is most likely to beg the question in favor of the point being made, especially if we distinguish between the convictions according to whether their fulfillment would cause harm to another.

<sup>22</sup> Moreover, should it turn out that there are reasons that can justify an act of killing but not an act of letting die when other things are equal, then this would support the main claim that negative and positive duties differ with respect to the reasons they admit as grounds for justified or excused violations. It would make it difficult, however, to capture this difference in the notion of "stricter."

<sup>23</sup> See Michael Wreen, "Breathing a Little Life into a Distinction", *Philosophical Studies*, vol. 46, 1984.

<sup>24</sup> See Allan Buchanan, "Justice and Charity", *Ethics*, vol. 97, no. 3. April 1987, pp. 558—575, especially section III.

<sup>25</sup> See Raymond Belliotti, "Negative Duties, Positive Duties, and Rights," *Southern Journal of Philosophy*, vol. 16, 1978.

<sup>26</sup> See Ronald Dworkin's account of the "trumping" effect of rights in *Taking Rights Seriously*, Harvard University Press, Cambridge, Mass. 1977. pp. 184—205.

<sup>27</sup> I am indebted to Ron Milo who, without complaint, read and commented on more drafts of this article than I care to mention.

*Department of Philosophy,  
Loyola University of Chicago,  
6525 N. Sheridan Road,  
Chicago, IL 60626,  
U.S.A.*