

## THE MORAL STATUS OF ENABLING HARM

Samuel C. Rickless

[Draft: Please do not cite without permission]

There is a puzzle at the heart of non-consequentialism that has not yet received a completely satisfying solution. The puzzle concerns the moral status of what Philippa Foot has described as “enabling” harm, i.e., withdrawing an obstacle that would, if left in place, prevent a pre-existing causal sequence from leading to foreseen harm. Non-consequentialists who take the distinction between doing and allowing harm to be morally relevant are divided about the proper moral classification of enabling harm. For some, *all* cases of enabling harm are morally equivalent to cases of *allowing* harm.<sup>1</sup> For others, at least *some* cases of enabling harm are morally equivalent to cases of *doing* harm.<sup>2</sup> The latter claim has also been endorsed by consequentialists, who treat it as a point in favor of their approach to ethics that there appears to be no principled non-consequentialist way of explicating the moral relevance of enabling cases.<sup>3</sup> The purpose of this paper is to clear up the existing confused state of non-consequentialist theorizing on this topic, and settle the debate by vindicating Foot’s original and far-seeing proposal.<sup>4</sup>

## THE DOCTRINE OF DOING AND ALLOWING

In a set of closely related publications, Philippa Foot describes what one might call “the doctrine of doing and allowing” (or “DDA,” for short).<sup>5</sup> According to the DDA, the distinction roughly captured in ordinary discourse by the terms “doing harm” and “allowing harm” is morally relevant in that the duty not to do harm is stronger than the

duty not to allow it to occur. Paradigm cases of doing harm, on Foot's view, consist in initiating or sustaining a causal sequence that leads to (foreseen) harm. To initiate a causal sequence is to set it in motion; to sustain it is to keep it going when it would otherwise have stopped. Thus, for example, if I set a trolley in motion in such a way that the trolley crushes a person who is trapped on the trolley tracks, then I have done harm by *initiating* a causal sequence that leads to it. And if I depress the trolley's accelerator pedal as the trolley is slowing down in such a way that the additional impetus results in the crushing of the trapped person, then I have done harm by *sustaining* a causal sequence that leads to it. Paradigm cases of allowing harm, on the other hand, consist in failing to prevent a pre-existing causal sequence from bringing about (foreseen) harm. (Henceforth, for the sake of convenience, I will drop the reference to the harm's being foreseen.) Thus, for example, if I could divert a speeding trolley that is bearing down on one trapped person onto an empty side-track but fail to do so, then I allow harm to befall the trapped person by failing to prevent the trolley from crushing him.

To say, as Foot does, that the duty not to do harm is stronger than the duty not to allow it to occur is to say, among other things, that when all else is equal, in situations in which one is forced to choose between doing harm and allowing harm, one is morally required to allow harm unless the harm that is allowed is out of all proportion to the harm that would otherwise be done. In particular, if one is forced to choose between doing harm to one and allowing the same kind of harm to befall five, one is morally required to avoid doing harm. A case of this kind is *Rescue*:

### *Rescue*

Roger is driving across a mountainous region to save five innocent people who are being slowly crushed to death by a large pneumatic press. If Roger reaches the press, he will be able to shut it off and the five will walk away unharmed. As Roger reaches a narrow mountain pass, he discovers that there is one innocent person trapped on the road ahead.<sup>6</sup>

In *Rescue*, Roger is forced to choose between two options. The first option is to stop the car, in which case the one lives and the five die. The second option is to drive across the mountain pass just in time to shut off the pneumatic press, in which case the one dies and the five live. It seems clear that it is morally impermissible for Roger to choose the second option over the first: he may not crush the one on his way to saving the five from being crushed. This kind of forced choice case brings out in the clearest possible way the inconsistency between the DDA and consequentialism, a theory that bids us to produce, from among all available alternatives, the one that leads to the best consequences overall. At the same time, if the only alternative to Roger's crushing the one is the failure to save a sufficiently large or disproportionate number (say, one million), then it may, on Foot's account, be permissible for Roger to crush the one. The duty not to do harm is therefore more stringent than the duty not to allow harm to occur, but it is not absolute.

### **ENABLING HARM**

Within a non-consequentialist theory that includes the DDA, what is the moral status of enabling harm? One view, championed by Foot herself, is that harmful enablings are

morally on a par with harmful allowings. Foot offers us the case of respirator removal as evidence for the claim that the removal of an obstacle that stands in the way of harm is morally equivalent to the failure to prevent some pre-existing causal sequence from leading to harm. (Call this claim “the Equivalence Hypothesis.) And, at least at first blush, this seems exactly right. One way to test the Equivalence Hypothesis is by considering cases in which one is forced to choose between allowing harm and enabling harm where the only additional putatively morally relevant difference between the two options is that the consequences of allowing are worse than the consequences of enabling. Here is one such case:

*Hospital*

In a hospital, a doctor has just plugged one person into the only available respirator. If the doctor either moves the one or unplugs him from the respirator, he will die. Five persons then arrive at the hospital and can be saved if and only if they are all plugged into the respirator at once. As it happens, the five will not survive being moved to the room with the respirator. But the respirator is movable.

In *Hospital*, the doctor is forced to choose between two options. The first option is to leave the one plugged into the respirator, in which case the one lives and the five die. The second option is to unplug the one from the respirator and plug the five into the respirator, in which case the one dies and the five live. (I ignore the clearly suboptimal option that involves unplugging the respirator from the one and then failing to attach it

the five.) Under the circumstances, it seems clear that it is at least morally permissible (and perhaps even morally required) for the doctor to choose the second option over the first. The Equivalence Hypothesis provides a straightforward explanation of this result. On the first option, the doctor allows the five to die in that she fails to prevent whatever is ailing them from causing their deaths. On the second option, the doctor enables the one to die in that she removes an obstacle (the respirator) that is preventing whatever is ailing the one from causing his death. So the doctor is forced to choose between allowing five to die and enabling one to die. If enabling harm is morally equivalent to allowing harm, then the only morally relevant difference between the options is that choosing the first will lead to greater harm than choosing the second. Given that it is morally desirable to minimize harm when all else is equal, it follows that it is morally permissible (and perhaps even morally required) for the doctor to choose the second option over the first.

If the Equivalence Hypothesis is true, then in situations in which an agent is forced to choose between doing harm to one and enabling the same kind of harm to befall five, the DDA predicts that it is morally impermissible to choose the former option. For the DDA says that it is morally impermissible to do harm to one even if the only alternative is to allow five to suffer the same kind of harm. But, according to the Equivalence Hypothesis, there is no morally relevant difference between allowing harm and enabling the same kind of harm. It follows directly from the conjunction of the DDA and the Equivalence Hypothesis that it is morally impermissible to do harm to one even if the only alternative is to enable five to suffer the same kind of harm. Is this prediction borne out? Indeed it is. For consider the following case:

### *Motion Detector*

George is driving on a motion sensitive road. At the same time, not far from where George is driving, a large boulder is rolling down a hill towards five innocent people who are trapped in the boulder's path. Luckily, there is a steel fence between the boulder and the five trapped people. If the boulder hits the fence, the five will be saved. However, George discovers that if he stops the car, the motion detector will set off an explosion that will destroy the fence before the boulder reaches it. In order to save the five, George resolves to keep driving. Unfortunately, George now discovers an additional innocent person trapped on the road ahead. If George keeps driving, he will crush and kill the one.

In *Motion Detector*, George is forced to choose between two options. The first option is to stop the car, in which case the one lives and the five die. The second option is to keep driving, in which case the one dies and the five live. Under the circumstances, it seems clear that it is morally impermissible for George to choose the second option over the first: he may not kill the one on his way to saving the five. In all morally relevant respects, then, *Motion Detector* is similar to *Rescue*, and this provides further confirmation for both the DDA and the Equivalence Hypothesis.

## **CHALLENGES TO THE EQUIVALENCE HYPOTHESIS**

The straightforward classification of harmful enablings as morally on a par with harmful allowings, which is part and parcel of Foot's original description and defense of the

DDA, has recently come under attack. Various theorists, both consequentialists and non-consequentialists, have argued that there are clear counterexamples to the Equivalence Hypothesis. Indeed, so they claim, there are many harmful *enablings* that are morally on a par with harmful *doings*. For Foot's non-consequentialist opponents, this is a sign that the DDA requires significant revision; for Foot's consequentialist opponents, it is a sign that the DDA is unprincipled and should be abandoned.

### *McMahan*

Let us begin by looking at Jeff McMahan's criticisms of the Equivalence Hypothesis. McMahan's challenge begins with a case he calls "Respirator":

#### *Respirator*

A person is stricken with an ailment that would normally be fatal but is given mechanical life-support to sustain him until the condition can be cured. While the patient is on a respirator, his enemy surreptitiously enters the hospital and turns the machine off. The patient dies. (1993, 254)

As McMahan sees it, if the distinction between enabling harm and doing harm "marks an intuitively morally important difference, then...we must conclude that it misclassifies this case" for "it is more natural to describe [*Respirator*] as a case of killing [that is, as a form of doing harm]; and we certainly evaluate it as such" (1993, 254). McMahan also offers the following potentially worrisome case:

### *Burning Building*

A person trapped atop a high building that is on fire leaps off. Seeing this, a firefighter quickly stations a self-standing net underneath and then dashes off to assist with other work. The imperiled person's enemy is, however, also present and, seeing his opportunity, swiftly removes the net so that the person hits the ground and dies.

In both *Respirator* and *Burning Building* the relevant victim's enemy removes an obstacle that would otherwise prevent harm (in both cases, death) from befalling the victim. So both cases are instances of enabling harm. But McMahan's intuition in *Burning Building*, no less than in *Respirator*, is that enemy's enabling victim's death is a killing that should be treated as morally on a par with other killings (such as driving over the one on the second option in *Rescue*.)

Although I share McMahan's intuitions about both cases, I would argue that they provide no evidence for the falsity of the Equivalence Hypothesis. The problem in both cases is that the agent whose actions enable the victim's death is described as the victim's *enemy*. Persons who are described as enemies of those whose death they bring about intentionally are generally understood to be acting with malicious intent. So, when faced with both cases, we take it for granted that enemy acts with malice aforethought. Our intuitions being responsive to what we take for granted, they reflect our evaluation of *malicious* harmful enablings, rather than our evaluation of harmful enablings *tout court*. Our intuitions therefore count as evidence for the claim that *malicious* harmful enablings

are morally on a par with *malicious* harmful doings, but they do not suggest that *non-malicious* harmful enablings are morally equivalent to *non-malicious* harmful doings.

The problem with both of McMahan's initial counterexamples to the Equivalence Hypothesis is reminiscent of the problem with James Rachels' well-known argument for the moral equivalence of doing and allowing.<sup>7</sup> Rachels' argument relies on a famous pair of cases. In the first case, Smith maliciously drowns his nephew to obtain an inheritance. In the second case, Jones maliciously fails to save his drowning nephew to obtain an inheritance. Rachels' intuition is that Smith's action is as morally reprehensible as Jones's inaction, and hence that there is no morally relevant difference between doing harm and allowing harm. But here too the most that could be said to follow from this pair of cases is that *malicious* harmful doings are morally on a par with *malicious* harmful allowings.

In response to this reply, McMahan might argue that his purpose in describing *Respirator* and *Burning Building* is not to establish that harmful enablings can be *morally* on a par with harmful doings. Rather, he might argue, his purpose is to do no more than "contribute to the identification of the empirical criteria governing the use of the concepts of killing and letting die," and perhaps more broadly, the concepts of doing harm and allowing harm (1993, 250). If this is correct, then the *point* of the two cases is to show that ordinary language classifies at least some enablings as killings, and more generally as harmful *doings*.

If this is indeed the point of *Respirator* and *Burning Building*, then I wholeheartedly concur with McMahan's conclusions about the application conditions of 'killing' and 'doing harm' in ordinary language. It seems perfectly appropriate to classify

enemy's action in *Respirator* as a killing, and as a case of doing harm; similarly for enemy's action in *Burning Building*. But, *from the moral point of view*, which is, after all, what is driving McMahan's investigation into the *concepts* of killing and letting die, these conclusions are neither here nor there. Foot's main thesis is *not* a claim about the ordinary language application conditions of terms like 'killing' and 'letting die'; it is, rather, a claim about the moral significance of a *technical* distinction (between allowings and enablings on the one hand, and doings and sustainings on the other) that, as she perhaps incorrectly surmised, maps roughly onto the ordinary language distinction between 'allowing' and 'doing'. What matters morally is not whether (some) enablings are properly described as 'doings' *in the ordinary sense*, but whether (some) enablings are to be treated as morally on a par with 'doings' *in the technical sense*. And no investigation into the application conditions of terms in ordinary language is capable of shedding light on this question.

Let us then return to the moral question. Are there any *non-malicious* harmful enablings that we think should be treated as morally equivalent to *non-malicious* harmful doings? McMahan, as it happens, considers a case of non-malicious harmful enabling, namely the following variant of *Burning Building*:

*Burning Building 2*

A person trapped atop a high building that is on fire leaps off. Seeing this, a firefighter quickly stations a self-standing net underneath. But he then immediately notices that two other persons have jumped from a window several

yards away. He therefore repositions the net so that it catches the two. The first jumper then hits the ground and dies. (1993, 262)

About this case, McMahan writes that “it seems absurd to say that the firefighter kills the one; rather he merely allows him to die” (1993, 262). That’s as may be. The moral question, however, is whether *Burning Building 2* counts as evidence for, or perhaps even evidence against, the Equivalence Hypothesis. And, interestingly, the answer is that the case provides striking *support for* the hypothesis. In *Burning Building 2*, once he has placed the net underneath the first jumper, the firefighter faces a choice between two options. The first option is to leave the net where it is, in which case the two jumpers die and the first jumper lives. The second option is to reposition the net underneath the pair of jumpers, in which case the two jumpers live and the first jumper dies. (I ignore the firefighter’s third and clearly worst option, which is to move the net so that it is underneath *none* of the jumpers, in which case all three jumpers die.) On the first option, the firefighter saves one but *allows* two to die. On the second option, he saves two but *enables* one to die. He will therefore save the greater number if he chooses option two. We may suppose that the firefighter’s choice is not guided by any malicious intentions; he is focused exclusively on helping. If non-malicious harmful enabling were morally on a par with harmful doing, then, according to the DDA it would be morally impermissible for the firefighter to reposition the net to save the greater number. And yet, intuitively, the firefighter is at least morally permitted, and perhaps even morally required, to reposition the net. So *Burning Building 2* strongly suggests that there is no morally

relevant difference between harmful enablings and harmful allowings. In this respect it is strikingly similar to *Hospital*.

The fact that McMahan interprets *Burning Building 2* as a case of letting die, rather than a case of killing, suggests that he shares my moral intuitions about the case. But McMahan also claims that in all variants of the original *Burning Building* case in which the harmful enabling is performed by someone who did not place the net under the first jumper (or who did not occupy the same role as the person who placed the net under the first jumper), the relevant enabling counts as a killing. And this suggests that he might balk at the claim that it is morally permissible to enable harm to save the greater number in these variants. But this would be a mistake. For consider the following variant:

*Burning Building 3*

A person trapped atop a high building that is on fire leaps off. Seeing this, a firefighter quickly stations a self-standing net underneath. He then quickly leaves the scene to assist those of his comrades who are battling the blaze. At the same time, Sally, a passer-by, then immediately notices that two other persons have jumped from a window several yards away. She therefore repositions the net so that it catches the two. The first jumper hits the grounds and dies.<sup>8</sup>

In this case, Sally faces the same options that the firefighter faces in *Burning Building 2*; she can allow one to be saved and allow two to die, or she can save the two and enable one to die. And it seems no less morally permissible for her to choose the second option

in this case than it is for the firefighter to choose the second option in *Burning Building 2*. It may be that Sally would ordinarily be described as having *killed* the one if she chooses option two (though I myself do not share this linguistic intuition), but again what matters is not whether her enabling is properly described as a ‘doing’ *in the ordinary sense*, but rather whether her enabling is to be treated as morally on a par with a ‘doing’ *in the technical sense*. And the intuitive answer to this question is clearly in the negative.

McMahan’s theory of the factors that determine whether we will ordinarily count a fatal enabling as a ‘killing’ or as a ‘letting die’ is complex and fascinating. One of his claims is that it is relevant to the evaluation of a fatal enabling as a killing or as a letting die whether the obstacle that is removed is “complete and self-sustaining, requiring no further contribution from [the one who originally provided it] to keep the threat at bay” (1993, 256). As evidence of this, McMahan offers us the following pair of cases:

*The Pipe Sealer*

An earthquake cracks a pipe at a factory, releasing poisonous chemicals into the water supply. Before a dangerous amount is released, a worker seals the pipe. But a year later he returns and removes the seal. As a result, numerous people die from drinking contaminated water. (1993, 256)

*The Dutch Boy*

A little Dutch boy, seeing that the dike is beginning to crack, valiantly sticks his finger in the crack to prevent the dike from breaking and flooding the town. He waits patiently but after many hours no one has come along who can help.

Eventually succumbing to boredom and hunger, the boy withdraws his finger and leaves. Within minutes the dike bursts and a flood engulfs the town, killing many. (1993, 257)

McMahan claims that our ordinary language intuitions bid us classify *The Pipe Sealer* as a case of killing but *The Dutch Boy* as a case of letting die. Again, that's as may be. The moral question, however, is whether the pipe sealer's removal of the seal is morally on a par with a harmful doing, and whether the Dutch boy's removal of his finger is morally on a par with a harmful allowing. And the answer here, it seems to me, is that both forms of enabling are, yet again, *morally* equivalent to harmful allowings, rather than to harmful doings. To see this, consider the following variants of both cases:

*The Two Pipes*

An earthquake cracks a pipe at a factory, releasing poisonous chemicals into a town's water supply. Before a dangerous amount is released, a worker seals the pipe. The worker realizes that if the seal is removed, one thousand people will die from drinking contaminated water. A little while later the worker returns to the factory to inspect the pipes. He notices a crack in a second pipe, which, if left unsealed, will release chemicals that will kill two thousand people in the next town over. Unfortunately, the only way to seal the crack in the second pipe is to remove the seal on the first pipe and transfer it to the second.

*The Dutch Boy and the Trolley*

A little Dutch boy, seeing that the dike is beginning to break, valiantly sticks his finger in the crack to prevent the dike from breaking, flooding the town, and killing one thousand people. He then notices a runaway trolley with two thousand people in it, hurtling down a nearby hillside, close to going over the edge of a cliff and killing everyone in it. The boy realizes that if he pulls a lever a few yards away, the trolley will stop before it goes over the edge. But he knows that in order to reach the lever he will need to remove his finger from the crack, and that if he does so the dike will break.

In *The Two Pipes* the worker has two options. He can leave the seal on the first pipe, thereby saving one thousand people but *allowing* two thousand to die; alternatively, he can remove the seal and place it on the second pipe, thereby saving two thousand people but *enabling* one thousand to die. (Again I ignore the worker's third and clearly worst option, which is to remove the seal from the first pipe and fail to place it on the second pipe, thereby enabling one thousand to die and allowing two thousand to die.) In *The Dutch Boy and the Trolley* the boy also has two options. He can leave his finger in the dike, thereby saving one thousand people but *allowing* two thousand to die; alternatively, he can remove his finger from the dike and pull the lever to stop the runaway trolley, thereby saving two thousand people but *enabling* one thousand to die. (I also ignore the boy's third and worst option, which is to remove his finger from the dike and fail to pull the lever.) If the harmful enablings in these cases were morally equivalent to harmful doings, then according to the DDA it would be morally impermissible for the worker to

remove the seal from the first pipe and place it on the second, and it would also be morally impermissible for the Dutch boy to remove his finger from the dike and pull the lever. And yet, intuitively, these are actions that the worker and the boy are at least morally permitted, and perhaps even morally required, to perform.

The upshot is that McMahan's purported counterexamples establish neither that enabling harm *per se* is morally inequivalent to allowing harm *per se*, nor that enabling harm *per se* is morally equivalent to doing harm *per se*. Some of McMahan's examples show at best that *malicious* harmful enablings are morally equivalent to *malicious* harmful doings. The rest of his examples show at best that particular cases of enabling harm can, under certain circumstances, fall under the ordinary language concept of doing harm. The moral is that McMahan has given us no good reason to abandon Foot's original claim that there is no morally significant difference between harmful enabling and harmful allowing.

## 2. *Vihvelin and Tomkow*

However, McMahan is not the only non-consequentialist opponent of the Equivalence Hypothesis. Recently, Vihvelin and Tomkow (henceforth, "V&T") have argued that some cases of enabling harm are morally on a par with doing harm. As their paradigms of doing harm and allowing harm, V&T use Bennett's examples, *Push* and *Stayback*:

### *Push*

A cart stands at the top of a hill. Agent pushes it. The cart rolls down the hill and fatally injures a child.

*Stayback*

The cart is already rolling; Agent could but does not interpose a rock which would stop it. The cart rolls down the hill and fatally injures a child. (Bennett 1995, 67; Vihvelin and Tomkow 2005, 192)

V&T then consider the following examples of enabling harm, the first of which is Bennett's:

*Kick*

The cart is rolling towards a point where there is a rock that would bring it to a halt. Agent kicks away the rock; the cart rolls down the hill and fatally injures a child.

*Dislodge*

A cart stands at the top of a hill, its wheel chocked by a rock. Agent kicks the rock away. The cart rolls down the hill and fatally injures a child. (Bennett 1995, 67; Vihvelin and Tomkow 2005, 193)

V&T describe *Kick* as a case of preventer prevention (for in kicking away the rock, Agent prevents the rock from preventing the child's death), and, following Schaffer (2000), describe *Dislodge* as a case of "causation by disconnection". Regarding these two cases of enabling harm, V&T argue as follows:

[W]hatever temptation there may be to use the language of ‘allowing’ in [*Kick* and *Dislodge*] there is no corresponding temptation to treat these agents *morally* on a par with the agent of *Stayback*. In real life cases of disconnection and preventer prevention—the hangman springing the gallows’ trap, the vandal who removes the guard rail at the tourist lookout or the greedy pharmacist who waters the vaccine—we count the agents as killers on a par with *Push* and not bystanders like *Stayback*.

V&T conclude that at least *some* fatally harmful enablings, though perhaps not all, are *morally* equivalent to killings, and, more generally, that at least *some* harmful enablings are *morally* equivalent to harmful doings.

But it should now be easy to see why V&T’s argument does not establish this conclusion. V&T claim that Agent’s kicks in *Kick* and *Dislodge* are actions that “we count as” killings. But, again, whether we *count* them as killings in ordinary language is irrelevant to the question of whether they are morally equivalent to harmful doings in the technical sense. What, then, supports V&T’s claim that Agent’s kicks in *Kick* and *Dislodge* are *morally* inequivalent to harmful allowings? The answer is that, on their view, these kicks are analogous to other cases of disconnection and preventer prevention, each of which is morally on a par with harmful doing. However, when we look at these cases more carefully, we can see that they do not establish what V&T think they do.

V&T mention three cases. Let us begin with the last two, “the vandal who removes the guard rail at the tourist lookout” and “the greedy pharmacist who waters the

vaccine”. Now, ordinarily, vandals are understood to be persons who have little or no regard for public or private property. Moreover, any person who removes a guard rail at a lookout knows, or should know, that the action could prove fatal to tourists who stop by. So removing the guard rail is grossly negligent at best, and malicious at worst. Further, a greedy pharmacist who waters the vaccine (presumably without informing the people who might be inoculated with it) is willing to trade the lives of innocent persons for his own benefit, and he too is guilty of gross malfeasance at best, malicious wrongdoing at worst. The fact that the intentions of the relevant agents in these two cases are not directed at the good serves as a distortion in the eliciting of intuitions about the moral status of harmful enablings *per se*. In this respect, these cases are similar to McMahan’s *Respirator* and *Burning Building*, cases in which the relevant agent is described as the “enemy” of the relevant victim. We can no more legitimately infer from the case of the vandal and the case of the greedy pharmacist that certain sorts of harmful enablings are morally equivalent to harmful doings than we can infer the same conclusion from *Respirator* and *Burning Building*. The most we can legitimately infer from the vandal case and the greedy pharmacist case is that *negligent, malfeasant, or malicious* harmful enablings are morally equivalent to *negligent, malfeasant, or malicious* harmful doings. And this is not sufficient to establish what V&T want to show.

But what of V&T’s first case, the case of “the hangman springing the gallows’ trap”? We may suppose, for argument’s sake, that the hangman has excellent moral reasons for springing the trap. Perhaps the person hanged was himself an unrepentant mass murderer who freely confessed and asked for death when he was captured. Isn’t *this* scenario a case of harmful enabling that is morally equivalent to harmful doing?

After all, in springing the trap, the hangman simply removes an obstacle to a causal sequence with a fatal upshot (namely, the sequence that begins with the hanged person's fall into the void, and ends with the tightening of the noose around his neck and consequent death by asphyxiation). And isn't the hangman's action morally equivalent to asphyxiating the confessed murderer directly?

Perhaps so. But again, the assumption that the hanged person is a confessed murderer, and hence deserves to be hanged, counts as a distorting element in the eliciting of moral intuitions about case. When someone deserves death, it may not matter morally whether his death is enabled or brought about directly. But it does not follow that there is no morally relevant difference between harmful enablings *per se* and harmful doings *per se*. To see whether there is such a morally relevant difference, it helps to change the case to one in which the person who is in danger of being hanged is innocent of wrongdoing and finds himself with a noose around his neck through no fault of his own. Then, to see whether springing the trapdoor is morally equivalent to initiating or sustaining a causal sequence that leads to his death, we can alter the circumstances further by adding a forced choice between springing the door and allowing a greater number of people to die in the same way. Here is such a case:

### *The Two Trap Doors*

One innocent person finds himself standing on a trap door with a noose around his neck. At the same time, five innocent people with nooses tied to long ropes around their necks are falling towards a second heavily padded, but open trap door. The two trap doors are mechanically connected so that if the first trap door

is sprung, the second automatically closes. Roger, a bystander, happens upon the lever that will spring the first trap door. Roger knows that if he pulls the lever, the first trap door will open and one innocent person will be hanged, while at the same time the second heavily padded trap door will close, and five innocent people will be saved. On the other hand, Roger knows that if he doesn't pull the lever, then the second trap door will remain open and five innocent people will be hanged, while at the same time the first trap door will remain closed and one innocent person will be saved.

Roger has a choice between enabling one to die while saving five or allowing five to die while allowing one to be saved. Intuitively, it seems morally permissible for Roger to pull the lever and save the greater number. But if harmful enablings were morally equivalent to harmful doings, then the DDA would speak against Roger's pulling the lever. The fact that we judge it morally permissible (perhaps even morally required) for Roger to pull the lever strongly suggests that harmful enablings are morally equivalent, not to harmful doings, but rather to harmful allowings.<sup>9</sup>

Another case that is sometimes used to support the claim that some harmful enablings are morally equivalent to harmful doings is the following scenario:

*Drive Away*

Suppose A's car is parked between a rolling rock and a helpless child. A sees the child in the rock's path, but drives off to avoid an ugly dent. (Boorse and Sorensen, 1988, 127; Vihvelin and Tomkow, 2005, 202-203.)

Christopher Boorse and Roy A. Sorensen (henceforth, B&S) claim that “A can be convicted of murder or manslaughter in any American court,” and conclude that there is no morally significant difference between this kind of enabling and a doing that leads to the same result (Boorse and Sorensen, 1988, 127). V&T (2005, 203) claim that “[t]he driver is as much as killer as the agent in *Push*, which is why the juries would convict.” But, again, from the moral point of view, whether we are inclined to describe the case as a killing is neither here nor there. It is, of course, morally reprehensible for A to drive away to avoid a dent. But it is also morally reprehensible for A (in a slightly different scenario) to fail to interpose his car between the rolling rock and the child if his reason for failing to do so is to avoid an ugly dent. And in this sort of case A *allows* the child to die. So the fact that A’s action is morally reprehensible in *Drive Away* does not show that harmful enablings are morally inequivalent to harmful allowings.

What matters is the question of moral equivalence, and here again it is easier to answer this question by altering the case into one involving a forced choice between enabling one to die and saving five or allowing five to die and saving one. What we need is a case much like *Burning Building 2*. Let’s call it *Drive Away 2*:

*Drive Away 2*

A’s car is parked between a rolling rock and one helpless child. A now sees a second rock rolling towards a group of five other helpless children. If A leaves his car where it is, the one helpless child will be saved but the five helpless children will die. But A can drive a short distance and stop, thereby interposing

his car between the second rolling rock and the five helpless children. If A does this, the five helpless children will be saved but the one helpless child will die.

In *Drive Away 2*, it seems morally permissible for A to drive the short distance and stop his car to save the five. Yet in doing so he enables one helpless child to die. So if harmful enabling were morally equivalent to harmful doing, then the DDA would speak against his moving the car. (Again, I ignore A's third option, which is to drive away from the scene, thereby enabling one child to die and allowing five more children to die.) *Drive Away 2* therefore suggests that harmful enablings are morally on a par with harmful allowings, rather than to harmful doings.

### 3. *Guns*

Thus far, we have considered and rejected a number of putative counterexamples to the Equivalence Hypothesis. But there is one particularly recalcitrant type of counterexample that some cite as evidence for the claim that at least some harmful enablings are morally equivalent to harmful doings, and hence morally inequivalent to harmful allowings. I now want to explain why this putative counterexample too does not work.

The example I have in mind is the kind of case in which one person shoots another. Frances Howard-Snyder considers the following scenario:

Sassan shot Victor. He pulled the trigger. The gun fired. A bullet flew out of the barrel and entered Victor's body. Victor died from the bullet wound. A clearer case of killing is impossible to find. (Howard-Snyder, 2007, section 6)

According to Howard-Snyder, it is *obvious* that Sassan kills Victor by shooting him. And indeed it does seem obvious. However, as Schaffer (2000, 287) points out, pulling the trigger of a gun involves “disconnecting the sear, allowing the spring to uncoil (propelling the striker onto the powder, compression of which produces the explosion which propels the bullet).” Unless the trigger is pulled, the sear prevents the explosion that leads to the firing of the bullet. So pulling a gun's trigger is a kind of enabling, for it involves the removal of an obstacle that lies in the way of an explosion that propels the fatal bullet. So when Sassan shoots Victor, his killing of Victor is an enabling. Moreover, it seems obvious, not only that Sassan *kills* Victor, but also that Sassan's killing of Victor is morally equivalent to *doing* harm. For there seems to be no morally relevant difference between Sassan's shooting Victor on the one hand and Sassan's pulverizing Victor on the other.

However, not even the Sassan case falsifies the Equivalence Hypothesis. First, as we saw earlier, the fact that Sassan kills Victor does not entail that harmful enablings are morally on a par with harmful doings. Because “doing harm” is a technical concept, it does not follow from the fact that something is correctly described in ordinary language as a ‘killing’ that it counts as a ‘doing’. Nor does the fact that many or most killings are doings entail that all killings are doings. So there is no inconsistency in saying both that Sassan kills Victor and that he does not *do* harm (in the technical sense) in killing Victor.

Second, it is easy to explain why there mistakenly appears to be no morally relevant difference between Sassan's shooting Victor and Sassan's pulverizing Victor. In both cases, it is obvious that Sassan does something wrong, indeed something morally reprehensible. It may therefore seem to follow that the type "fatal shooting" is morally indistinguishable from the type "fatal pulverizing". But it does not follow. This is the familiar point brought out by the failure of Rachels' Smith-Jones pair of cases. Just as the fact that Jones's allowing his cousin to drown is just as morally reprehensible as Smith's drowning his cousin does not entail that allowing harm is morally equivalent to doing harm, so the fact that Sassan's shooting Victor is just as morally reprehensible as Sassan's pulverizing Victor does not entail that enabling harm is morally equivalent to doing harm. The confounding parameter in both pairs of cases is the same, namely the malicious intent of the relevant agent of harm.

To correctly gauge the moral status of fatal shootings, it helps to consider another forced-choice scenario, one in which the agent must choose between enabling one to die while saving five and allowing five to die while one is saved. Here is one such case:

*The Gun and the Bomb*

A bomb is set to go off in five minutes. Five innocent people are strapped to the bomb, and all five will die if the bomb goes off. Several yards away, there is a gun encased in concrete, pointed at the heart of one innocent person, also encased in concrete some distance away. If the gun is fired, the one innocent person will die. Reina, a savvy bystander who has just arrived on the scene, notices that the gun's trigger is linked to the bomb's detonator by a motion-detector: if the trigger

remains motionless, then the detonator will fire and the bomb will explode; but if the trigger is pulled, thereby allowing the propulsion of the striker onto the powder and the consequent explosion that propels the bullet towards the one, then the detonator will be deactivated and the bomb rendered harmless.

Reina has a choice. On the one hand, she can leave the trigger alone, thereby making sure that one lives but *allowing* five people to die. On the other hand, she can pull the trigger, thereby saving five people but *enabling* one person to die. If pulling the trigger were morally equivalent to a fatal doing, then the DDA would speak against her doing so. But intuitively it is at least morally permissible for Reina to pull the trigger. It follows that pulling the trigger is not morally equivalent to a fatal doing, but rather to a fatal allowing.

And yet the intuition that shooting someone is morally on a par with doing harm is stubborn. Why so? There is a simple explanation. Most of us are ignorant of how guns actually work. Moreover, it is easy to imagine that a gun's trigger functions as a kind of detonator, even though what the trigger does is hold a sear that prevents a spring from uncoiling. I suggest that those of us who readily think of firing a gun as morally equivalent to a doing think that firing a gun *is* a kind of doing, rather than a kind of enabling. Because of this, it is likely that nothing but careful attention to the inner workings of a gun will lead ordinary folk to change their minds about the moral status of gun firings generally. By contrast, it is not difficult to elicit the judgment (in, say, *Burning Building 2*) that removing a safety net is morally equivalent to a harmful allowing, rather than to a harmful doing. The reason, I suggest, is that there are no

prejudices about how safety nets function to confound judgment. Notice also that most *actual* harmful shootings in our experience are morally on a par with *actual* harmful doings. It is rather difficult to think of a *realistic* case in which the pulling of a gun's trigger is morally distinguishable from a harmful doing: cases like *The Gun and the Bomb* are not exactly a dime a dozen. So even when the folk learn how guns work, it can be difficult to get them to see that pulling the trigger is morally equivalent to a harmful allowing. This is because they are mistakenly inclined to generalize from the numerous instances in which shooting a human being is just as morally reprehensible as pulverizing him.

### CONCLUSION

In recent years, it has become increasingly common to assume, on the basis of simple (non-forced-choice) cases of disconnection and preventer prevention, that some harmful enablings are morally equivalent to harmful doings. Indeed, several authors have built (or have suggested) their own rather complex versions of the DDA on this assumption.<sup>10</sup> As I have argued, none of these cases succeeds in disproving Foot's original claim that the moral status of enabling harm is equivalent to the moral status of allowing harm. We are left to conclude that there is no morally relevant difference between harmful enablings and harmful allowings *per se*, and hence that the more complex versions of the DDA that are based on the rejection of this conclusion are false.

## REFERENCES

- Bennett, Jonathan. 1995. *The Act Itself*. Oxford: Clarendon Press.
- Boorse, Christopher, and Roy A. Sorensen. 1988. "Ducking Harm." *Journal of Philosophy* 85: 115-134.
- Foot, Philippa. 1967. "The Problem of Abortion and the Doctrine of Double Effect." *Oxford Review* 5: 5-15. Reprinted in *Virtues and Vices and Other Essays in Moral Philosophy* (Oxford: Blackwell, 1978), pp. 19-32, and in Steinbock and Norcross (1994), pp. 266-279.
- Foot, Philippa. 1984. "Killing and Letting Die." In *Abortion: Moral and Legal Perspectives*, edited by Jay L. Garfield and Patricia Hennessey (Amherst: University of Massachusetts Press), pp. 177-183. Reprinted in Steinbock and Norcross (1994), pp. 280-289.
- Foot, Philippa. 1985. "Morality, Action, and Outcome." In *Morality and Objectivity: A Tribute to J. L. Mackie*, edited by Ted Honderich (London: Routledge & Kegan Paul), pp. 23-38.
- Hanser, Matthew. 1999. "Killing, Letting Die and Preventing People from Being Saved." *Utilitas* 11: 277-295.

Howard-Snyder, Frances. 2002. "Doing vs. Allowing Harm." *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, URL = <http://plato.stanford.edu/entries/doing-allowing/>.

Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Clarendon Press.

McMahan, Jeff. 1993. "Killing, Letting Die, and Withdrawing Aid." *Ethics* 103: 250-279.

Rachels, James. 1975. "Active and Passive Euthanasia." *New England Journal of Medicine* 292: 78-80. Reprinted in Steinbock and Norcross (1994), pp. 112-119.

Schaffer, Jonathan. 2000. "Causation by Disconnection." *Philosophy of Science* 67: 285-300.

Steinbock, Bonnie, and Alastair Norcross (ed.). 1994. *Killing and Letting Die*, Second Edition. New York: Fordham University Press.

Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. Oxford: Oxford University Press.

Vihvelin, Kadri, and Terrance Tomkow. 2005. "The Dif." *Journal of Philosophy* 102: 183-205.

Woollard, Fiona. 2008. "Doing and Allowing, Threats and Sequences." *Pacific Philosophical Quarterly* 89: 261-277.

## NOTES

---

<sup>1</sup> See Foot (1967) and Hanser (1999).

<sup>2</sup> See Boorse and Sorensen (1988), McMahan (1993), Vihvelin and Tomkow (2005), and Woollard (2008).

<sup>3</sup> See Rachels (1975), Kagan (1989), Bennett (1995), Unger (1996), and Howard-Snyder (2007).

<sup>4</sup> There are other puzzles for non-consequentialism, including puzzles concerning the proper rationale for treating the distinction between doing harm and allowing harm as morally significant, and puzzles concerning the causal metaphysics behind the distinction (for example, whether there is a metaphysical basis for distinguishing between enabling harm and sustaining harm). These puzzles require special treatment, and I do not propose to solve them here. In previous work (see -----), I have explained why, as Foot also argued, the distinction between doing and allowing does not collapse into the distinction between act and omission.

<sup>5</sup> See Foot (1967; 1984; 1985).

<sup>6</sup> This case is inspired by the famous Rescue II case described by Foot (1984).

---

<sup>7</sup> See Rachels (1975).

<sup>8</sup> This case resembles the case of *Two Cars and a Rock* described by Hanser (1999, 285).

<sup>9</sup> We can imagine a similar reformulation of the vandal case. Imagine that a bystander happens upon a lever that is connected to two guard rails, one of which is up and the other one of which is down. As it happens, a runaway trolley with one innocent person in it is heading towards the first guard rail (the one that is up), while a runaway trolley with five innocent people in it is heading towards the second guard rail (the one that is down). If the bystander pulls the lever, the first guard rail will go down while the second guard rail will go up. If the bystander pulls the lever, he enables one to die but five are saved from certain death. If the bystander does not pull the lever, he allows five to die but one person is saved from certain death. It seems morally permissible for the bystander to pull the lever. Again, for similar reasons to the ones pointed out in the text, this case strongly suggests that harmful enablings are morally far closer to harmful allowings than they are to harmful doings.

<sup>10</sup> See, for example, McMahan (1993), Vihvelin and Tomkow (2005), and Woollard (2008).