Are Enabling and Allowing Harm Morally Equivalent?

KASPER LIPPERT-RASMUSSEN
University of Aarhus

It is sometimes asserted that enabling harm is morally equivalent to allowing harm (the moral equivalence thesis). In this article, I criticize this view. Positively, I show that cases involving self-defence and cases involving people acting on the basis of a reasonable belief to the effect that certain obstacles to harm will remain in place, or will be put in place, show that enabling harm is harder to justify than allowing it. Negatively, I argue that certain cases offered in defence of the moral equivalence thesis fail, because either (1) their similarity with the archetypal trolley case limits their relevance to an assessment of this thesis, or (2) they are compromised by their reliance on the elusive notion of a situation being completely stable.

I. INTRODUCTION

In her classic article 'The Problem of Abortion and the Doctrine of the Double Effect' Philippa Foot distinguished between two kinds of allowing:

> There is firstly the allowing which is forbearing to prevent. For this we need a sequence thought of as somehow already in train, and something that an agent could do to intervene . . . So, for instance, he could warn someone, but allows him to walk into a trap. [There is secondly the allowing: KLR] which is roughly equivalent to enabling; the root idea being the removal of some obstacle which is, as it were, holding back a train of events. So someone may remove a plug and allow water to flow; open a door and allow an animal to get out; or give someone money and allow him to get back on his feet.[1]

Both kinds of conduct contrast with doing harm, paradigm cases of which consist in 'initiating or sustaining a harmful causal sequence'. In what follows I shall say that the second kind of conduct 'enables', while the first merely 'allows', an outcome.[2] My question is: where

[1] P. Foot, *Virtues and Vices and Other Essays in Moral Philosophy* (Berkeley and Los Angeles, 1977), p. 26. The distinction leaves us with some difficult cases. Suppose a potential killer and an unloaded gun are in a room. Rather than removing the gun, I load it and put it back where it was. Arguably, this does not fall into Foot's first category, because there is no sequence in train which will lead to the shooting of the killer's victim, and which I decline to intervene in by loading the gun. Nor, however, does the case fall clearly into Foot's second category: I do not simply remove an obstacle – the gun not being loaded – that is holding back a train of events resulting in the killing.

[2] S. C. Rickless, 'The Moral Status of Enabling Harm', *Pacific Philosophical Quarterly* 92 (2011), pp. 66–86, at 66, characterizes enabling harm as 'withdrawing an obstacle that would, if left in place, prevent a pre-existing causal sequence from leading to foreseen

© Cambridge University Press 2015
Utilitas Vol. 27, No. 3, September 2015
doi:10.1017/S0953820815000035

CrossMark
https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

the outcome is harm to another, are these cases of non-doing morally equivalent?

Foot thinks that, on some occasions, allowing harm is 'just as bad as' doing harm. However, on other occasions it is clear that doing harm is worse than allowing harm, and on no occasions is allowing harm worse than doing harm. *Ceteris paribus*, then, doing harm is morally worse than allowing harm. Because this claim is not my focus in this article, and because common-sense moral intuitions clearly support it, I shall take it for granted.

This summary of Foot's views points to two questions about enabling harm: is it as bad as doing harm? Is it as bad as allowing harm? The simplest responses here either demote enabling to the moral status of doing (both being worse than allowing) or promote enabling to the moral status of allowing (both being better than doing). A more complicated position locates enabling midway: it is less bad than doing harm but worse than allowing harm. An even more complicated position would be that enabling harm is sometimes worse, sometimes less bad, than doing harm, and similarly so with allowing harm – it has no fixed, comparative moral status of this sort per se.

No philosopher who has addressed these questions has opted for the most complicated answer.[3] Some – e.g. Vihvelin and Tomkow[4] – have defended the claim that enabling harm is, sometimes at least, as bad as doing harm, while others have argued that enabling harm is intermediate between doing and allowing harm.[5] Yet, others have claimed that enabling harm is (i) morally on a par with allowing harm, and (ii) less bad than doing harm.[6] I shall refer to (i) as the *moral equivalence thesis*. This thesis is the focus of my article. It implies that cases of allowing harm and enabling harm have the same moral status (e.g. one cannot be morally permissible when the other is not)

harm'. On this approach one cannot enable unforeseen harms (nor allow them). This epistemic clause is not part of Foot's characterization of enabling harm. Foot's use of the terms 'thought of' and 'as it were' in the quoted passage suggests that she takes series of events that allow and enable outcomes to differ, not in respect of intrinsic features of the events, but in how we think of them. I set aside this interpretation of Rickless's view.

[3] In this article, I simply criticize the moral equivalence thesis. Its falsity is compatible with one of the more complicated answers, so for all I say here those answers remain on the table.

[4] K. Vihvelin and T. Tomkow, 'The Dif', *Journal of Philosophy* 102 (2005), pp. 183–205, at 194.

[5] C. Barry, G. Øverland, 'The Feasible Alternatives Thesis: Kicking Away the Livelihoods of the Global Poor', *Politics, Philosophy and Economics* 11 (2012), pp. 97–119; C. Barry, M. Lindauer and G. Øverland, 'Doing, Allowing, and Enabling Harm', *Oxford Studies in Experimental Philosophy*, ed. J. Knobe, T. Lombrozo and S. Nichols (Oxford, 2014), pp. 62–90.

[6] M. Hanser, 'Killing, Letting Die and Preventing People from Being Saved', *Utilitas* 11 (1999), pp. 277–95; Rickless, 'The Moral Status'.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

when all other factors, such as the badness of their consequences, are held constant across the relevant cases. For instance, the moral equivalence thesis implies that allowing and enabling harm, even if they are not intrinsically morally different, do not interact differently with other moral factors, such as beliefs.[7] This means that, as long as all other morally relevant factors are constant, a case of allowing harm in which the agent has a certain belief about the consequences of his conduct and a case of enabling harm in which the agent has the same belief about the consequences of his conduct *must* have the same moral status.

The main conclusion of this article is that the moral equivalence thesis is false. Section II presents a counterexample to the conclusion of an argument given by Samuel Rickless in favour of the thesis. The counterexample involves self-defence. Also, it raises a doubt about Rickless's premise that enabling harm is morally on a par with allowing harm in the example he offers. Section III considers variants of the case Rickless's argument builds on. It argues that the moral significance of allowing and enabling harm can differ in situations where an agent acts on reasonably formed expectations about how to avoid harm. Section IV shows that certain cases of enabling harm, where it is permissible to enable one person to be harmed to save others, are best explained, not by the truth of the moral equivalence thesis, but by the similarities between these cases and standard trolley scenarios where it is even permissible to harm one to save more. Section V rebuts the way in which friends of the moral equivalence thesis have tried to accommodate recalcitrant moral intuitions by appealing to the distinction between a stable situation and one where causal sequences in train are set to continue. Section VI concludes.

Before we turn to my critique of the moral equivalence thesis, there are three preliminaries. First, the equivalence thesis is significant for the two reasons expounded in this and the following paragraph. The moral inequivalence of doing and allowing harm has become something of a hallmark of deontological thought. This is so, even if, conceptually speaking, deontologists and non-deontologists need not disagree over this issue.[8] Within deontology there is pressure therefore to explain why doing and allowing differ. By far the most plausible explanation

[7] For the notion of interaction between morally relevant factors, see J. Dancy, *Moral Reasons* (Oxford, 1993); S. Kagan, 'The Additive Fallacy', *Ethics* 99 (1988), pp. 5–31; F. M. Kamm, 'Killing, Letting Die: Methodological and Substantive Issues, *Pacific Philosophical Quarterly* 64 (1983), pp. 297–312; J. McMahan, 'Killing, Letting Die, and Withdrawing Aid', *Ethics* 103 (1993), pp. 250–79, at 274).

[8] Some deontologists think that doing and allowing (as opposed to, say, intending and merely foreseeing) harm are morally equivalent, and that in principle a consequentialist can affirm that doing harm has greater negative value than allowing harm.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

attaches moral significance to the distinction between what the victim of harm has independently of the person who countenances harm (either by doing, allowing, or for that matter enabling harm) and what the victim of harm has only through the person who countenances harm. By way of illustration, consider Judith Jarvis Thomson's famous case of the violinist in need of life support who has been hooked up with another person without this person's consent.[9] On Thomson's view, the supporter would kill the violinist if she were to disconnect him, yet this would be morally permissible. Frances Kamm suggests that the explanation of killing being permissible in this case is that it is an untypical killing in that the violinist would only lose something that she has through the killer/supporter, i.e. continued life. As Kamm puts it: the baseline relative to which one must not harm the violinist is 'the condition he would have been in if he had never been attached to his supporter'.[10] This motivates the following argument: because countenancing harm is less wrong when the person harmed loses only what he has through the agent of harm – as Thomson's case illustrates – and because it is true by definition in cases of allowing harm that the victim loses only what he could have had through the agent who does not assist, but not similarly true by definition in cases of killing, it follows that allowing harm is less wrong per se than doing harm.[11] However, if enabling and allowing harm are morally equivalent, this argument for the moral significance of the distinction between doing and allowing harm fails. This is so, because in cases of enabling harm, where the enabler of harm is different from the one who provides the obstacle to harm if such an agent exists, the victim is deprived of a good she would have had if the enabler of harm had not been present at the scene. Given this, it seems that in order to offer a plausible rationale for deontology one would have to defeat the moral equivalence thesis.[12] Hence, the present issue is theoretically important.

A further reason why the equivalence thesis is significant relates not so much to moral theory as to practical, real-life moral concerns. Thomas Pogge has influentially argued that citizens of rich countries do not merely allow global poverty. Rather, they bring it about through

[9] J. J. Thomson, 'A Defense of Abortion', *Philosophy and Public Affairs* 1 (1971), pp. 47–66, at 47–9.

[10] F. M. Kamm, *Bioethical Prescriptions* (Oxford, 2013), p. 187.

[11] For an account along these lines, see Kamm, 'Killing'. For critical discussion, see S. Kagan, *The Limits of Morality* (Oxford, 1989).

[12] It should be noted, however, that in his arguments against the moral equivalence thesis Rickless seems to presuppose that the distinction between doing and allowing has moral significance. Accordingly, he cannot (nor does he) employ the moral equivalence thesis to undermine the moral asymmetry between doing and allowing harm.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

the imposition of an international regime skewed in their favour at the expense of the global poor, and accordingly they have much stronger moral obligations to eliminate poverty than it is normally assumed. In Pogge's view we must 'stop thinking about world poverty in terms of helping the poor'.[13] In response, Barry and Øverland have submitted that once we substitute the tripartite distinction between doing/allowing/enabling harm for the bipartite distinction between doing/allowing harm with which Pogge operates we see that, typically, citizens in rich countries enable the global poor to be harmed (e.g. through the erection of trade barriers) rather than doing harm to them (e.g. through military interventions). Although they reject the moral equivalence thesis, Barry and Øverland allow that it 'would undermine quite significantly Pogge's argument that affluent people have particularly stringent moral reasons to address global poverty'.[14] They also hold that rejection of the thesis does not rescue Pogge's position. They take enabling harm to be a distinct moral category intermediate between doing and allowing harm; and in their view 'moral reasons associated with contributing to harm through enabling it seem to be less constraining than reasons associated with doing harm' though more constraining than those associated with allowing harm.[15] On this basis, with certain additional assumptions, they conclude that Pogge exaggerates the stringency of the duties of affluent citizens to eliminate global poverty.[16]

Second, I have formulated the moral equivalence thesis in comparative terms that require us to see actions as better (less bad) or worse (more bad). However, the bald assertion that one act is worse than another is bound to convey rather different things to different people. One person might detect a Kantian theme: conduct deteriorates, morally, as it fails to track duty and/or is done from motives other than the sense of duty. Another person will simply assume that an action of the worse kind has the following feature: it can be morally impermissible when an otherwise comparable action of the less bad kind is not.[17] It is this latter sense of badness that I and, so I believe, other people who have written on the topic have in mind. For instance, Rickless appears to have something like this sense in mind when he

[13] T. Pogge, *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms* (New York, 2008), p. 23.

[14] Barry and Øverland, 'The Feasible', p. 112.

[15] Barry and Øverland, 'The Feasible', p. 110.

[16] Barry and Øverland, 'The Feasible', p. 113.

[17] This is compatible with the possibility that the acts differ contingently in other, permissibility-relevant, dimensions in a way licensing us to say that, in those dimensions, the worse action is better than the less bad action.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

writes: 'it is more difficult to justify doing harm than it is to justify allowing harm'.[18]

Third, it is interesting that some enablings of lethal harm are best viewed as killings, while others are naturally described as cases of 'allowing someone to die'. What criteria guide us here is a very good question,[19] but not one that I can address in this article.

II. RESCUING AND SELF-DEFENCE

I start by considering an important argument recently given in support of the moral equivalence thesis by Samuel Rickless. Rickless sketches a case in which enabling and allowing harm seem to be morally alike – in fact, the case is a modification of one of a series of cases developed by Jeff McMahan in the course of arguing that at least some cases of lethal withdrawal of aid are morally equivalent to killing (hence the '3'):[20]

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.[21]

According to Rickless, Sally acts permissibly. He claims this supports the moral equivalence thesis.[22] For while we might ordinarily describe the case as one in which Sally kills the single individual, it is a case of her enabling harm to the single person trapped in the building in Foot's sense. After all, Sally simply removes an obstacle to the process in which the lone individual falls to his death. The largely uncontroversial fact that she can permissibly enable harm to one person to save two (each) from the same harm suggests that the enabling is morally equivalent to allowing harm. It does so because, as participants in the debate acknowledge, one is permitted to allow harm to one person to save two from it. If enabling harm were morally equivalent to doing harm, however, one would expect to conclude that Sally acted

[18] Rickless, 'The Moral Status', p. 66; see also T. Scanlon, *Moral Dimensions* (Cambridge, 2008), pp. 3, 102–5.

[19] See J. Bennett, *The Act Itself* (Oxford, 1995), pp. 62–73; Barry et al., 'Doing'.

[20] McMahan, 'Killing', pp. 263–4.

[21] Rickless, 'The Moral Status', p. 73.

[22] We can set aside John Taurek-like reasons for doubting the permissibility of allowing harm to a few to save the many based on the notion that each person should be given a fair chance (whatever that is) of being saved. See J. Taurek, 'Should the Numbers Count?', *Philosophy and Public Affairs* 6 (1977), pp. 293–316.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

impermissibly. Normally it is impermissible to do harm to one person to save two from it. This, I take it, is Rickless's argument.[23]

There are at least two reasons to be sceptical about this argument, one of which is internal to Rickless's own example and one of which is not. The first is this. It is doubtful whether, in Rickless's example, enabling harm really is morally equivalent in all respects to allowing the same harm. If the single victim could somehow paralyse Sally before Sally repositions the net below a second burning building, where there is also one innocent person trapped, this might be permissible. In any case, intuitively, his doing this would be less wrong than the act would be in a slightly different case in which he paralyses Sally, who is at that moment considering whether to position the net below the first or the second building, and this somehow results in the net being placed below the first building.[24] In this case, he would clearly act impermissibly.[25] This does not immediately contradict anything Rickless says, but it suggests that it might not be enough, for the purpose of testing the moral equivalence thesis against the case of Burning Building 3, to look merely at the permissibility of repositioning of the net.[26]

[23] Obviously, the argument is not deductively valid, nor does Rickless claim that it is. It strikes me as a plausible inference to the best explanation, though, for reasons set out below, ultimately I think it is unwarranted.

[24] This is true even if we compare cases where Sally has promised not to enable/allow harm to the respective jumpers – more harm can permissibly be imposed on Sally by the recipient of the promise to avoid the bad consequences of Sally's breaking her promise by enabling harm than that which can permissibly be imposed on Sally by the recipient of her promise, i.e. to the relevant jumper, to avoid the bad consequences of her breaking her promise not to allow harm. In Burning Building 3 and the variations of it which I consider Sally has not issued any such promises to the persons involved; nor has she otherwise communicated to them that they can rely on her acting in a certain way.

[25] Cf. McMahan, 'Killing', p. 255. The findings of Barry et al., 'Doing', suggest that most people take agents who knowingly enable harm to have stronger compensatory duties to victims than agents who knowingly allow harm. (Strictly speaking, they use various test cases where the agent is innocent because she is unaware of the harm she countenances. Arguably, this is not the best way to test the significance of innocence, because one can be culpably negligent in bringing about a harm one was not aware that one might bring about.) On the assumption that, generally speaking, there is a connection between the stringency of compensatory duties to victims and the amount of harm a potential victim can impose on a potential harm-doer to avoid harm, these findings provide some support for my claim about our relevant moral intuitions here. (Admittedly, Barry et al., 'Doing', p. 83, think that the compensatory duties of agents who unknowingly enable harm are less stringent than the compensatory duties of agents who unknowingly do harm. Even if this is right, and right because compensatory duties of agents who unknowingly enable or allow harm are morally on a par, this does not rescue the moral equivalence thesis. Enabling and allowing harm would still interact differently with knowledge of the harm one countenances in relation to duties of compensation, and thus in this way they are morally inequivalent.)

[26] Cf. Barry et al., 'Doing', p. 82. This suggestion is similar to Kamm's observation that in James Rachels's pair of examples of the uncle countenancing his nephew's drowning (which Rickless appeals to) one is entitled to harm the uncle to prevent the nephew from

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

Second, this suspicion is strengthened when we consider simpler cases of self-defence not involving a third party trying to maximize the number of people saved. Suppose you and I are both trapped in unowned cages. A hungry bear is outside punching the cage doors. Eventually, it will succeed in opening one of them, eat the person inside, and then, no longer being hungry, depart leaving the other person unharmed. In the first case to consider I can reinforce the door to my cage, thereby allowing the other person to be eaten. In a second case, I can weaken the door of the other person's cage with my metal saw, thereby enabling the other person to be eaten. Intuitively, my conduct in the first case is less wrong than it is in the second.[27] As Barry, Lindauer and Øverland put it: 'reasons against enabling harm appear to constrain conduct more with regard to shifting costs from ourselves to others than do reasons against allowing harm'.[28] Indeed, all things considered, my conduct in the first cage case seems morally permissible, unlike my conduct in the second case. Hence, further support for the case against the moral equivalence thesis can be derived from cases of self-defence.[29] I conclude that Rickless's argument appeals to a false premise and has a conclusion which, on other grounds, we have reason to think is false. Below I shall offer further reasons for scepticism about both it and the moral equivalence thesis.

III. EXPECTATIONS AND THE ALLOWING/ENABLING HARM DISTINCTION

In this section I will argue, *pace* the moral equivalence thesis, that allowing and enabling harm interact differently with conduct based on expectations about harm. To do this I shall consider some variants of Burning Building 3.

Consider first the *no-enabling harm belief case*. In this case the person trapped at the top of the blazing building acts on the reasonable belief that a self-standing net has been positioned next to his building. That is, unlike his counterpart in Rickless's Burning Building 3, he jumps *after* forming the reasonable belief that the firefighter has placed the net under the building and, as Sally is aware, he has no reason to suppose that Sally might move it. Compare this with the *no-allowing harm belief case*. Here the firefighter unfolds the net, but then dashes

drowning in the doing-case in ways that do not translate to the allowing-case. See J. Rachels, 'Active and Passive Euthanasia', *New England Journal of Medicine* 292 (1975), pp. 78–80.

[27] Cf. C. Boorse and R. Sorenson, 'Ducking Harm', *Journal of Philosophy* 115 (1988), pp. 115–34, at 115–17.

[28] Barry et al., 'Doing', p. 65.

[29] I am grateful to Susanne Burri for input at this point.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

off to attend to something else. The trapped man notices Sally passing by and, noting that Sally has picked up the net, reasonably forms the belief that Sally will position the net under his window. He then leaps, believing that the net will save his life. Again, we can suppose Sally is aware that he holds this belief.[30]

Intuitively, these cases differ. The fact, in the *no-allowing harm belief case*, that the trapped man reasonably believes Sally will not *allow* him to die and acts on the basis of this belief clearly does not render it impermissible for Sally to reposition the net when doing so will save two people. By contrast, the fact, in the *no-enabling harm belief case*, that the trapped man reasonably believes he will save his life by jumping, and that Sally will not *enable* him to be harmed, and acts on these beliefs, makes it intuitively impermissible for Sally to reposition the net. Sally would not be permitted to remove the net if it were only one person jumping from the other window. Also, it would be permissible for her to allow, say, the wind to move the net to the other window. In this variant of Burning Building 3, then, the enabling of harm does not appear to be morally equivalent to its being allowed – the two sorts of conduct interact differently with reasonably formed beliefs upon which the man threatened with harm acts.[31]

It might be replied that, even if this is so, it remains true that in the basic Burning Building 3 case Rickless presents repositioning the net as permissible. However, in order to defeat the moral equivalence thesis I need not deny – even if, as indicated in the previous section, I think there is a case for denying – that enabling harm is equivalent to allowing harm in cases where the agents involved have not acted on the assumption that a certain obstacle will be in place preventing them from being harmed. All I need claim is that the two are morally inequivalent in cases where the agents involved have acted on such an assumption.

[30] It makes no important moral difference to my assessment of this pair of cases whether Sally is aware of the jumper's beliefs. However, an anonymous reviewer found the example more convincing with this stipulation. Moreover, an experiment conducted by Barry et al., 'Doing', shows that in cases of enabling harm many do think it makes a moral difference to compensatory obligations whether the agent of enabling harm is aware of the consequences of what she is doing (see n. 26). Because I only need to identify one case where allowing and enabling harm interact differently with expectations in order to defeat the moral equivalence thesis, I have no reason not to make this stipulation.

[31] The claim that enabling and allowing harm interact differently with reasonably formed beliefs upon which persons threatened with harm act is compatible with the following claim: that an agent has a duty to make some effort to avoid causing others to form and act upon reasonable but false beliefs about what he will allow to occur. Presumably, such a duty is less stringent than the similar duty to prevent others from reasonably forming, and then acting upon, false beliefs about what he, the agent, will enable.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

Rickless's appeal to Burning Building 3 could be buttressed, alternatively, by saying the following: my argument draws illegitimate support from the contingent truth that, generally, people in a position to be saved by an observed obstacle to harm that is already in place expect to be saved, while people who might be saved by an obstacle to harm if a person, whom they can see, intervenes and repositions the obstacle do not expect to be saved. It could be said that it is, or at least feels, morally worse to countenance harm to people who do not expect it than to people who do. Our sense that removing the net to save two is impermissible in the *no-enabling harm belief case* might reflect this moral difference – we tacitly assume that the two people at the other window believe they are doomed, while the lone man thinks he is safe.

However accurate these points are, they cannot explain the moral difference between the *no-allowing* and the *no-enabling harm belief cases*. And in any case the built-in bias can be removed. Thus consider a slightly different case in which Sally acts as the enabler of harm: the *equalized no-enabling harm belief case*. Here the two people at the second window jump again on the basis of a reasonable expectation that Sally will reposition the net and, thus, not allow them to be harmed. They know that Sally has often been in such situations before, and that she has always sought to maximize the number of lives saved. Like the lone man, then, they now have a reasonable expectation that in jumping from the building they will be saved. If our intuition that Sally's removal of the net in the *no-enabling harm belief case* is impermissible was determined by our belief that, in cases where harm is not actively *done* to another, it is morally worse to allow/enable harm to be suffered by people who believe they are safe than it is to allow/enable harm to be suffered by people who think they are doomed, we should find that our intuition changes now that this factor has been neutralized. (Everyone now thinks it is safe to jump.) However, this does not happen – at least, in my intuitive judgement. Hence, reasonable expectations formed on the basis of an obstacle's actually having been positioned are different, morally speaking, from reasonable expectations formed on the basis of an agent's disposition to reposition obstacles. This suggests the following: enabling harm to a victim who acts on a reasonable expectation of not being harmed that was formed on the basis of a belief about an obstacle to harm having been positioned by someone other than the agent enabling the harm is worse than an otherwise comparable case of allowing harm.[32] Given this, we must conclude

[32] The moral difference between the two cases cannot be explained by the expectations of the two sets of jumpers because the jumpers do not differ in this respect: both expect to be rescued by the net if they jump. It might be replied that their expectations do differ in one respect: the lone jumper expects that Sally will not *enable* her to be harmed, while

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

that at least some forms of enabling harm are worse than otherwise comparable cases of allowing harm: allowing and enabling harm are not morally on a par if they interact differently with expectations. Hence the moral equivalence thesis is false.

I will now try to provide further support for this claim by looking at cases of burning buildings in which the persons in question face different options that each will give them some chances of survival and choose among them on the basis of their beliefs about the presence or non-presence of obstacles to harm:

Burning Building 4
A person is trapped atop a high building. He can either jump, which in the absence of a net located underneath will give him a 1 per cent chance of survival, or try to exit the building using the stairs, which will give him a 25 per cent chance of survival. Seeing this, a firefighter quickly stations a self-standing net below, which means a 95 per cent chance of survival for the person if the net remains in place, and he jumps. The firefighter then leaves the scene to assist those of his comrades who are battling the blaze. Seeing that the net has been placed below the building the person jumps, reasonably believing this to be, by now, his best chance of surviving. At the same time, Sally, a passer-by, then immediately notices that two other persons are in an identical predicament. They have jumped from a window several yards away acting on the reasonable belief, formed on the basis of Sally's behaviour on previous occasions, that Sally will reposition the net, and, like the lone man, they have not tried to exit the building via the stairs. Sally therefore repositions the net so that it catches the two. The first jumper hits the grounds and dies.

Intuitively, Sally acts impermissibly in this case. (Uncontroversially, she would do so in a one versus one case, and in such a case the single person might permissibly use force to prevent the repositioning of the net.) In fact, it seems to me that even if Sally were another firefighter, not a passer-by, it would be impermissible for her to reposition the net.[33] Now compare Burning Building 4 to 5:

Burning Building 5
A person is trapped atop a high building. He can either jump, which in the absence of a net located underneath will give him a 1 per cent chance of survival, or try to exit the building using the stairs, which will give him a 25 per cent chance of survival. Believing that it is very probable that the firefighter will place a net below, he leaps off the building. Seeing him jump, a firefighter quickly stations a self-standing net underneath, which means a 95 per cent chance of survival for the jumping man if the net remains in place. The firefighter then leaves the scene to assist those of his comrades who

the pair of jumpers expect that Sally will not *allow* them to be harmed. However, if the difference between these two expectations is morally significant, the best explanation of it is that the difference between enabling and allowing harm is morally significant. From what else could the moral significance of the distinction derive?

[33] This applies even if the two people act on the reasonably formed belief that Sally would not allow them to die.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

are battling the blaze. At the same time, Sally, a passer-by, then immediately notices that two other persons in an identical predicament 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.

This case seems morally different from Burning Building 4. For here it is not clear that Sally acts impermissibly when she repositions the net. It makes a difference that in Burning Building 5 the lone man does not jump as a result of having formed a belief induced by the firefighter's positioning of the net. It is harder to justify enabling harm to such an agent than it is merely to justify allowing it to an otherwise comparable agent. What makes the two cases different, morally, is not whether the lone person jumped believing that the net would be placed under his building or jumped believing it had already been placed there. This is morally irrelevant. The crucial difference is that in Burning Building 4, in contrast to Burning Building 5, the firefighters act in a way that leads the lone person to assess his situation in a particular way, and then to act in a way that will make him even worse off if the agent enables harm than he would have been had the agent not intervened at all. By way of support for this claim consider a case where the presence of a coast guard ship has led some sailors sitting on an adjacent shipwreck to believe reasonably that they will be picked up if they jump into the cold water and swim the short distance to the ship. Once the men are in the water and, let us suppose, unable to get back on the shipwreck, the coast guard ought not to abandon them simply because he can rescue more sailors elsewhere. The shipwrecked sailors gave up on another way of saving their lives – staying put on the shipwreck – in response to the presence of the coast guard ship and the more favourable opportunity this led them to believe they had. I conclude that this is a second kind of case (the first was the case of self-defence mentioned at the end of section II) which indicates that allowing and enabling harm are morally inequivalent. In both kinds of case, allowing and enabling interact differently with the choices people make on the basis of reasonable beliefs about whether others will enable or allow one to be harmed.[34]

[34] Allowing and enabling harm also interact differently with malicious intentions. Rickless, 'The Moral Status', p. 71, accepts that our 'intuitions . . . count as evidence for the claim that *malicious* harmful enabling is morally equivalent to harmful doing'. However, given that malicious harmful allowing is not morally equivalent to harmful doing it follows, *pace* Rickless, that allowing and enabling harm are not morally equivalent. To see that malicious harmful allowing is not morally equivalent to harmful doing consider a variation of an example provided by Jeff McMahan which Rickless, 'The Moral Status', p. 70, discusses in this context. A person is trapped atop a burning high building. Fortunately, there is a self-standing net below. Unfortunately, the wind will blow it away before the person reaches the ground if he jumps. This person's enemy is present at the scene and could prevent the net from being removed by the wind, but refrains from so

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

IV. THE BURNING BUILDING: A TROLLEY-LIKE CASE

In this section I want to consider a fourth reason to be sceptical about Rickless's argument as it is presented in section II. This reason applies to several of the other cases that have been offered in favour of the moral equivalence thesis.

A possible source of error weakening Rickless's appeal to the intuition that repositioning the net in Burning Building 3 is permissible is this: almost everyone acknowledges that there are cases where one can permissibly kill one person in an attempt to save more people. One case of this sort is the famous trolley example, in the Thomson variant, where a bystander can divert a trolley from a track on which five innocent people are trapped onto another track where only one innocent person is also trapped.[35] Quite how we should interpret this acceptable harming, and specifically what makes it an exception to the wider moral prohibition on killing one person when that will save five (if such prohibition there is), is, of course, highly controversial. But we need not enter that controversy for present purposes. The important point is that, if, and to the extent that, our assessment of Burning Building 3 is influenced by factors of the sort that persuade us to view the killing permissible in the trolley case, then it is an error to take the intuition that it is permissible for Sally to reposition the net to support the moral equivalence thesis – at least, without further argument. We already know there are some cases in which killing one to save more is permissible (just as it is permissible to allow one to die when the alternative is to let even more people die). Given this, enabling to die could be morally equivalent to killing despite the permissibility of Sally's repositioning of the net in Burning Building 3. In other words, it may be that the permissibility of enabling harm in this case is due to the factors that explain why, in the trolley case, it is permissible to kill one to save more.[36]

doing out of malice. Here the trapped person could not, say, permissibly shoot his enemy if doing so would make the enemy collapse in front of the net, thereby preventing it from being blown away from the building. However, he could permissibly use lethal force in a comparable case of doing harm (e.g. to prevent someone from setting the building on fire thereby trapping him with no means of escape).

[35] J. J. Thomson, 'Killing, Letting Die and the Trolley Problem', *Rights, Restitution, and Risk: Essays in Moral Theory*, ed. W. Parent (Cambridge, MA, 1986), pp. 78–93 ; cf. Hanser, 'Killing', p. 286 n. 10. An exception is J. J. Thomson, 'Turning the Trolley', *Philosophy & Public Affairs* 36 (2008), pp. 359–74; see, however, W. J. FitzPatrick, 'Thomson's Turnabout on the Trolley', *Analysis* 69 (2009), pp. 636–43. However, while Thomson now thinks it is impermissible to divert the trolley thereby killing one and saving five, it seems that this judgement on her part is not a moral intuition, but a theory- or argument-generated view.

[36] For present purposes I need not identify what these features are. All I need to deny is (1) that the features are common to all cases of enabling harm and allowing harm, on the one hand, and some cases of doing harm, on the other; and (2) that this explains why

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

Rickless comes even closer to drawing on trolley-like intuitions. He offers a case (reminiscent of cases discussed by Boorse and Sorenson and by Vihvelin and Tomkow) in which, in his view, enabling harm is equivalent to allowing harm:[37]

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.[38]

With a similar argumentative aim, Hanser appeals to a similar example.[39]

Drive Away 2 looks like a circumstantially adjusted trolley case, inasmuch as all of the innocent children in it face the same kind of threat and the agent can save five innocents by killing one.[40] Admittedly, with the pair of rolling rocks it involves two distinct (but qualitatively identical) threats, but this does not look like a significant moral difference. Suppose all six children are trapped in a narrow canyon and the one child is closer to the rolling rock than the other five. A's car is parked in front of the one child, but if it remains there the rock will hit the car, bounce into the air, and, having passed the one child, crush the other five. A can park his car in front of the five instead, saving them with the side-effect that the one child is crushed. If A were to do so, this would be morally no different from the decision, made by A in Drive Away 2, to move his car. If this is right, it suggests that the special features of the trolley case do not include the fact that the case involves the redirection of just one threat as opposed to several qualitatively identical threats. Admittedly, some might deny that this is right, contending that it makes a moral difference whether we are

enabling and allowing harm are morally equivalent and morally less wrong that harming. In support of this denial (of (1), specifically) consider a case where I enable someone else to kill a third person by giving him a gun without which he would otherwise have been unable to kill. This is a case of enabling harm that is no less wrong, morally, in virtue of having trolley-like features.

[37] Boorse and Sorenson, 'Ducking', p. 127; Vihvelin and Tomkow, 'The Dif', p. 196.

[38] Rickless, 'The Moral Status', p. 78.

[39] Hanser, 'Killing', p. 285.

[40] Hence, I disagree with Rickless when he writes: 'So if harmful enabling were morally equivalent to harmful doing, then the DDA [i.e. the doctrine of doing and allowing: KLR] would speak against his moving his car. Drive Away 2 therefore suggests that harmful enabling is morally equivalent to harmful allowing, rather than to harmful doing' (Rickless, 'The Moral Status', p. 78). The conditional in the first sentence is false, because DDA does not apply, or at least does not apply in the right way, to trolley-like situations.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

dealing with the same or distinct threats. It is implausible to claim that the fact the agent is dealing with one and the same threat makes it harder to justify intervention, so the claim has to be that it makes it easier. But Rickless is in no position to push this line of argument, because he thinks intervening in Drive Away 2 *is* permissible. On this view, despite the fact that it involves two distinct threats, the intervention must be treated as sufficiently like a trolley case so far as permissibility is concerned.

At this point, one can go either of two ways. One might argue that, in the trolley case, the agent who diverts the trolley does not, in the relevant sense, kill the person on the single-person track, but only enables him to die. Despite the fact that diverting the trolley to save the greater number is normally seen as a case of killing, the positioning of the tracks can be seen as an obstacle standing between the lone person trapped on the track and the threat.[41] However, this suggestion strikes me as highly revisionist. Suppose a gun will eventually fire and kill someone, say, fifty years in the future, unless someone pulls the trigger. I do so and as a result someone is hit by the bullet fired from the gun and dies. We would resist describing this as a case of enabling harm, so we should also resist the idea that the trolley case involves enabling harm.

Alternatively, one could infer – and this is the general lesson that I want to draw from the present section – that the permissibility of countenancing harm to the one helpless child in Rickless's case is explained, not by the fact that the case is one of enabling harm, but rather by the situation's trolley-case character. Accordingly, and in contrast with what Rickless contends, the fact, in this case, that you may enable one to die when you thereby save five from dying does not support the moral equivalence thesis.

V. WHEN DOES ONE ENABLE HARM?

The last three sections criticized Rickless's positive arguments for the moral equivalence thesis. They also presented reasons – relating to self-defence, and to the interaction between the allowing/enabling distinction and expectations of harm, plans of actions based thereon, and intentions (see note 35) – to reject it. In this section I will assess Rickless's attempt to fortify the moral equivalence thesis by showing that he can accommodate a case which seems to amount to a decisive counterexample to it.

[41] E.g. F. Howard-Snyder, 'Doing vs. Allowing Harm', *Stanford Encyclopedia of Philosophy*, < http://www.plato.stanford.edu/entries/doing-allowing/> (2011), p. 2.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

When assessing the moral equivalence claim, or for that matter its denial, it is crucial to clarify the conditions under which one can be said to remove 'some obstacle which is, as it were, holding back a train of events' rather than initiating a new 'train of events'. Rickless acknowledges this need and although he does not 'attempt to provide a complete account of the metaphysics of causal sequences', he believes that certain otherwise intractable cases of enabling death that seem morally like killing can be dismissed as counterexamples to the moral equivalence thesis on the ground that, unlike cases that are similar to the standard trolley case, such as the Drive Away 2 case discussed in section IV, they do not involve any threatening 'pre-existing causal sequence'.[42] In this connection he considers a case in which a bystander can pull a lever. If the bystander does so a trap door will open under an innocent person with a noose around his neck. If he refrains from doing so, another trap door on which five innocent persons are standing with nooses around their necks will open. Intuitively, it is impermissible to pull the lever. Since the bystander is only enabling the innocent person to die,[43] the fact that it is impermissible to pull the lever suggests that 'at least one kind of harmful enabling is morally equivalent to harmful doing'.[44] Call this example The Hanging.

In response to this putative counterexample to the moral equivalence thesis, Rickless maintains that, despite what one may initially think, The Hanging is not a case of enabling harm: 'The opening of the trap door does not make it possible for a causal sequence that is already in train to continue, for there is no such causal sequence. In particular, the one is not in the process of falling. The situation is, rather, completely stable.'[45] Accordingly, the fact that it is impermissible to pull the lever does not even speak to, let alone support, the moral equivalence thesis.

For three reasons, Rickless's treatment of this putative counterexample to the moral equivalence thesis is unsatisfactory. First, without any explanation it employs a notion of enabling harm that deviates from Foot's own conception. When one removes a plug allowing water to flow, presumably, the situation prior to unplugging might have been 'completely stable' in Rickless's sense, with no water flowing. Yet this

[42] Rickless, 'The Moral Status', p. 79.

[43] See Foot's examples of enabling reported above in section I.

[44] Rickless, 'The Moral Status', p. 79; cf. Vihvelin and Tomkow, 'The Dif', p. 194.

[45] Rickless, 'The Moral Status', p. 80; see also p. 83. One would like to ask: 'Completely stable with regard to what?' The answer cannot be that it is completely stable with regard to the survival of the innocent person. There is no causal sequence already in train whose continuation will result in the victim's death in the absence of intervention by any agent. Moreover, if this were the answer the situation would also be completely stable in Burning Building 3.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

is one of the examples that she uses to illustrate her second notion of allowing. Admittedly, Foot's notion of enabling harm might be unusual here, since it might be more natural to describe the pulling of the trap-door releasing lever as a case of harming rather than simply enabling harm.

Second, it is unclear that complete stability of this kind is morally relevant. Suppose the innocent person happens to be in the process of falling: the trap door is opening by a centimetre a year, and before the innocent person's death from natural causes it will eventually open entirely. In this case there is a 'causal sequence that is already in train to continue'. Hence, on Rickless's account, we have a case of enabling harm. If this is so, and if the moral equivalence thesis is true, the act of pulling the lever should differ morally (other things being equal) from the lever-pulling in the case Rickless discusses. Yet there seems to be no such difference. In my view this speaks against equivalence.

It might be said that the mere fact that the trap door is opening by a centimetre a year until the innocent person is hanged is irrelevant. The crucial point is that there is no causal sequence in train that is being held back by some obstacle in a way that prevents the innocent person from being killed *in the immediate future*. The causal sequence that will lead to the victim's death thirty years later is irrelevant to a moral assessment of pulling the lever.

This way of defending the moral equivalence thesis arguably involves a double-edged sword. Consider McMahan's Dutch boy, who has stuck his finger in a crack in a dyke to prevent it from crumbling.[46] The boy need not do anything to keep his finger there, or so we shall suppose. It is stuck in the crack, and to withdraw it would involve effort and time.[47] Suppose that even with the finger in the crack the dyke is still crumbling, however, at a reduced pace. If at some point the Dutch boy decides to free his finger, the dyke will crumble much faster. Intuitively, in this case, in withdrawing his finger the Dutch boy enables the harm that ensues. Yet, if we say that the fact that the dyke is slowly crumbling is a different causal sequence from the one that will lead to the destruction of the dyke if the boy withdraws his finger –

[46] McMahan, 'Killing', p. 257.

[47] This is unlike McMahan's version of the case, where the Dutch boy seems to make continuous efforts to keep his finger in the crack. My version shows that if something counts as withdrawing, (a) the provider of the obstacle must *do* something in order to prevent the removal of the obstacle, and (b) it must *cost* the one who put it there in the first place something not to withdraw the obstacle. Note also that this case is one in which the situation is, in some sense, completely stable. No intervention by any agent is needed to prevent the dyke from cracking. Yet it does seem to be a case of enabling harm if the Dutch boy withdraws his finger.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

which we must, if we offer the reply sketched in the previous paragraph – we can no longer say the boy is enabling harm.[48]

Third, whether a situation presents as a 'completely stable' one depends on the way (or the level at which) it is described, and it is far from easy to see how a descriptive decision could reasonably determine moral permissibility. At one level, for instance, it might be quite correct to say that someone's standing on a mountain peak is a completely stable situation. Yet, at another level, it might be incorrect, because continents are moving, and mountains are being eroded, and so on.[49] Note also that there is a sense in which a situation in which a person is placed in a respirator is completely stable, although most of us believe that doctors simply enable a death to occur when they remove it.[50] For the three reasons given, I suggest that Rickless fails to explain away the counterexample to the moral equivalence thesis presented by The Hanging.

VI. CONCLUSION

This article has taken a closer look at some cases offered in response to the claim that enabling harm is morally equivalent to allowing harm. It has argued that these are not compelling. Cases of enabling harm to one or more people to save even more people from comparable harm are similar in important respects to the archetypal trolley case. As a result the permissibility of enabling harm seems to tell us little about the moral status of enabling harm relative to the moral equivalence thesis. It has also been argued that the countenancing of harm is harder to justify when people act on the basis of a reasonable belief to the effect that certain obstacles to harm will remain in place than it is when they act on the basis of a reasonable belief to the effect that others will not allow harm to them. Cases involving self-defence or malicious intentions also appear to rebut the moral equivalence thesis. Finally, I have raised some doubts about the moral relevance of a situation

[48] Consider a case where, to prevent the dyke from being washed away, the Dutch boy places his finger in the crack. This does not completely stop the destruction of the dyke, but it slows it down very considerably. Now suppose that after some time, the boy decides to withdraw his finger, and that in the process of doing so, inevitably, he will damage the dyke in such a way that the destruction is accelerated (relative to the situation before his intervention). Suppose that given the time he has already spent with his finger in the crack, and given the accelerated rate of destruction after his withdrawal, the dyke will break at exactly the same time it would have broken had he never intervened in the first place. Will it still count as enabling harm now if he withdraws his finger? Would it do so if he were to accelerate the rate of destruction so much that the dyke will break before it would have done had he never intervened in the first place?

[49] Bennett, *The Act*, p. 68.

[50] It might be replied that the situation cannot go on indefinitely. However, if this is the relevant criterion hardly any situation is completely stable.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press

being completely stable, and some doubts about the serviceability of the notion of stability as a determinant of whether an act counts as an enabling of harm. Twenty years ago, reflecting on the moral significance of the distinction between killing and letting die in the light of cases involving enabling death, Jeff McMahan conjectured that 'an ethical theory that aims to unify and systematize the intuitive elements of common moral thought will have to be quite complex'.[51] This article supports that conjecture to the extent, at any rate, of showing that common moral thinking on the enabling of harm is more complex than the moral equivalence thesis allows.[52]

lippert@ps.au.dk

[51] McMahan, 'Killing', p. 274.

[52] Previous versions of this article were presented at the University of Oslo, 19 June 2012; the University of Roskilde, 5 April 2013; and the Society for Applied Philosophy Conference, University of Zurich, 29 June 2013. I thank Christian Barry, Saba Bazargan, Susanne Burri, Tom Douglas, Matthew Hanser, Ingmar Persson, Thomas Søbirk Petersen, Rob Reich, Samuel Rickless, Jesper Ryberg, Alex Voorhoeve, Garrath Williams, Gerhard Øverland, and an anonymous reviewer of this journal for excellent comments and criticisms.

https://doi.org/10.1017/S0953820815000035 Published online by Cambridge University Press