

promoting access to White Rose research papers



Universities of Leeds, Sheffield and York
<http://eprints.whiterose.ac.uk/>

This is an author produced version of a paper accepted for publication in **Topics in Contemporary Philosophy: Action, Ethics and Responsibility**.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/3415/>

Frowe, H. (2007) *Killing John to save Mary: a defence of the moral distinction between killing and letting die*. In: Campbell, J., O'Rourke, M. and Silverstein, H. (eds.) *Topics in Contemporary Philosophy: Action, Ethics and Responsibility*. *Topics in Contemporary Philosophy* (7). MIT Press, Cambridge, Massachusetts. (In Press)

**“Killing John to Save Mary:
A Defence of the Moral Distinction Between Killing and Letting Die.”¹**

Helen Frowe

Introduction

This paper defends the moral significance of the distinction between killing and letting die.² In the first part of the paper, I consider and reject Michael Tooley’s argument that initiating a causal process is morally equivalent to refraining from interfering in that process. The second part disputes Tooley’s suggestion it is merely external factors that make killing *appear* to be worse than letting die, when in reality the distinction is morally neutral. Tooley is mistaken to claim that we are permitted to kill bystanders who had no fair chance to avoid being at risk of harm. We can support the significance of the killing / letting die distinction by considering the difference between what we are permitted to do in self-defence against those who are going to kill us, and what we can do against those who are going to let us die. I also suggest that we are less responsible for the deaths we allow than for the deaths that we cause, since we do not make people worse off for our presence in cases where we fail to save them.

I. Causing Versus Failing to Interfere

(i) The Moral Symmetry Principle

Michael Tooley argues that it makes no difference whether one initiates a causal process leading to an event E, or one merely fails to interfere in that same process to prevent E occurring. This, he claims, shows the absence of any moral difference between killing and letting die. Tooley calls this the ‘Moral Symmetry Principle’.

Let C be a causal process that normally leads to an outcome E. Let A be an action that initiates process C, and B be an action that stops process C before outcome E occurs.

¹ I am particularly grateful to Brad Hooker, Russ Payne and Andrew Williams for detailed comments on earlier drafts of this paper. I owe thanks to the two referees for helpful comments. Also to Peter Cave, Adrian Moore, Charlie Pelling, Harry Silverstein, Philip Stratton-Lake, Galen Strawson, Jussi Suikkanen, and Mark Young for comments on an earlier draft of this paper.

² James Rachels’s ‘Active and Passive Euthanasia’ (Norcross & Steinbock (Eds) *Killing and Letting Die*, Fordham University Press, New York (1994)) is a classic defence of the claim that there is no moral difference between killing and letting die. Jonathan Bennett similarly denies that there is any significant difference between doing and allowing in *The Act Itself* (OUP, New York (1995)). Consequentialist views in general deny that killing is *in itself* any worse than letting die (see, for example, Shelly Kagan *The Limits of Morality* (OUP, Oxford (1989))).

Assume further that actions A and B do not have any other morally significant consequences, and that E is the only part or outcome of C which is morally significant in itself. Then there is no moral difference between performing action A, and intentionally refraining from performing action B, assuming identical motivation in the two cases.³

So, we have action A, which leads to outcome E. Doing A will cause E. The performance of action B will prevent E occurring. Intentionally not doing B will count as letting E happen.

Tooley illustrates the Moral Symmetry Principle using *Poisoned Whiskey*.

Poisoned Whiskey: Two brothers decide (independently of each other) to ill their wealthy father. The first brother poisons his father's whiskey. The second brother catches the first in the act. The second brother had been planning to poison the whiskey himself. Now he merely fails to warn his father that the whiskey has been poisoned, and also refrains from giving him an antidote to the poison.

Tooley claims that the actions of the two brothers in *Poisoned Whiskey* are morally equivalent, thus demonstrating that the Moral Symmetry Principle is sound.

We have a comparison between intentionally initiating a causal process leading to the father's death and intentionally *letting* the father's death occur. In the context of *Poisoned Whiskey*, Tooley's analysis seems fairly compelling. It is hard to pick out a moral distinction between the different actions of the two brothers. Tooley concludes that *Poisoned Whiskey* demonstrates the insignificance of the killing / letting die distinction.

(ii) *The Traditional Objection*

However, Tooley does acknowledge that there is a traditional objection to denying the significance of the killing / letting die distinction. The distinction might not seem to matter in *Poisoned Whiskey*. But there are many cases in which we *do* think it matters

³ Michael Tooley 'An Irrelevant Consideration: Killing Versus Letting Die', in Norcross & Steinbock (Eds.) *Killing and Letting Die*, Fordham University Press, New York (1994) p. 104

whether an agent kills or lets die. A universal acceptance of the Moral Symmetry Principle, therefore, will commit one to equating actions that are intuitively different.

An objection along these traditional lines tends to take the form of *Enemy Torture*.

Enemy Torture: A captured spy refrains from divulging information to his enemy captor, despite the fact that the captor will torture a child as long as the spy says nothing.⁴

This seems relevantly analogous to *Poisoned Whiskey*. The second brother refrains from administering the antidote to his father, despite knowing that his father will be killed as long as he does nothing. The spy refrains from revealing his information, despite knowing that the child will be tortured as long as he says nothing. The antidote and the information are both ways of interfering to prevent some event occurring. Tooley equates the actions of the two brothers in *Poisoned Whiskey*. But we do not seem to want to equate the actions of the spy and the captor.

In anticipation of this sort of objection, Tooley offers the following clarification of the Moral Symmetry Principle. Properly understood, it will become clear that, “this example is just not relevant to the moral symmetry principle.”⁵ Tooley argues that,

The moral symmetry principle states, very roughly, that it is as wrong intentionally to refrain from interfering with a causal process leading to some morally significant result as it is to initiate the process. It does *not* assert that it is as wrong to refrain from *preventing someone else* from initiating a causal process as it is to initiate it oneself.⁶

The Moral Symmetry Principle would fall prey to counter-examples like *Enemy Torture* only if it were committed to equating a failure to stop *others* from initiating a causal process with initiating the causal process oneself. And, Tooley claims, this is *not* something to which the Moral Symmetry Principle is committed.

⁴ Tooley, p. 105

⁵ Tooley, p. 105

⁶ Tooley, p. 105

It matters for Tooley that in *Enemy Torture*, the spy fails to interfere to prevent the *captor* from *initiating* a causal process of torture. In *Poisoned Whiskey*, however, the second brother fails to interfere in an *existing* causal process. The idea seems to be that it is not as important to stop other people initiating causal processes as it is to stop existing causal processes (assuming that one is able to do so). So, the Moral Symmetry Principle equates failing to stop the existing process of poisoning by administering the antidote with poisoning the whiskey. But it does not equate failing to prevent the *initiation* of the poisoning with poisoning the whiskey. And nor does it equate failing to stop the captor initiating the torture of the child with torturing the child oneself. The Moral Symmetry Principle, it seems, can resist the traditional objection to denying a moral difference between killing and letting die.

This is what Tooley's position looks like:

(1) *Initiating the causal process*

is equivalent to:

(2) *Failing to interfere with the existing causal process.*

But (1) is *not* equivalent to:

(3) *Failing to prevent someone else from initiating the causal process.*

But we need only tweak *Enemy Torture* very slightly to show the problem with this position. Imagine that in *Revised Enemy Torture* the captor is *already* torturing the child, and will not stop unless the spy divulges his information. Saying that one does not have a responsibility to interfere with the *initiation* of a causal process is not going to resolve the problem posed by *Revised Enemy Torture*. Divulging the information in this case no longer counts as interfering with the *initiation* of a process, but as interfering with an *existing* process. Not divulging information once the captor has begun torturing the child will still come out as equivalent to torturing the child oneself.

To try to solve this problem, we need to know *why* Tooley thinks there is a difference between failing to prevent the continuation of an existing causal process, and failing to prevent the initiation of that process. Tooley suggests that there is a *prima facie* reason not to equate these two actions:

... the intuitive feeling of most people would surely be that the mere fact that when one prevents someone else from doing something one is interfering with someone's action, whereas when one merely refrains from doing something oneself one is not, is a morally relevant difference. Thus there is a *prima facie* reason against any extension of the moral symmetry principle that would have the consequence that intentionally refraining from preventing someone else from doing something is morally equivalent to doing it oneself.⁷

So, whether one fails to prevent the continuation or the initiation of a process is not in *itself* the important difference. This difference matters only because of some *further* fact about the initiation of causal processes. This further fact is that in order to prevent the initiation of a process one would have to interfere with another person's actions.

Note that the point at which Tooley requires the second brother to act in *Poisoned Whiskey* is *after* the first brother's actions have stopped. Poisoning is a self-sustaining process. Once he has added the poison to the whiskey, the first brother need do nothing else to bring about his father's death. Stopping the process, then, will not require the second brother to interfere in anyone else's actions. The second brother can either warn his father that the whiskey has been poisoned, or administer the antidote once the poison has been imbibed. He can do this without interfering with his brother's actions. His failure to do so, therefore, makes him as culpable as if he had added the poison himself.

The same is not true of *Revised Enemy Torture*. Torturing the child may well involve continued action on the part of the captor.⁸ The spy, we can imagine, would have to stop the captor from *doing* something, whereas the first brother in *Poisoned Whiskey* is no longer doing anything that the second would have to interfere with to stop his father's death. This, Tooley suggests, gives us a reason not to equate these actions. But is this a sufficient difference to mark out a moral distinction between a failure to interfere in *Poisoned Whiskey* and a failure to interfere in *Revised Enemy Torture*?

(iii) *Interfering With Other People's Actions*

⁷ Michael Tooley 'An Irrelevant Consideration: Killing Versus Letting Die', p. 105

⁸ Tooley can employ this distinction only if has broad, non-physical, conception of what counts as interfering with a person's actions, such that the spy's revealing of his information counts as such an interference.

Why would the fact that a particular part of a causal process involves another person's actions make a moral difference to our obligation to interfere with that process? How is Tooley going to motivate this principle? In an earlier paper, Tooley's formulation of the Moral Symmetry Principle specifies that interference in an existing causal process should involve only "a minimal expenditure of energy".⁹ And interfering with other people's actions does tend to involve a greater expenditure of energy than not doing so. However, this caveat is notably absent from the version of the principle that Tooley offers in the later article that I am discussing here. I think we should look at why this clause has been removed.

If Tooley specifies that refraining from saving is equivalent to killing only if saving would require very little effort on the part on the agent, this is going to substantially restrict the number of situations to which his principle can apply. And as the principle applies to fewer and fewer cases, we can ask legitimate questions about its plausibility. Why, if we can posit a morally significant distinction in so many other cases, would we deny the distinction's relevance in these limited cases to which Tooley's principle can apply?

I suspect that this is why Tooley has removed the clause about minimal expenditure of energy. There is little point in establishing a principle the application of which is strictly, and implausibly, limited. If Tooley wishes to maintain that interference in other people's actions really is a morally relevant difference, he will need some other rationale that will more easily generalise.

The only other candidate for this rationale seems to be something like a principle about interfering with other people's autonomy. But we do not generally regard this as a reason to stop a person from doing something very wrong. Whatever reasons the spy might have for failing to prevent the captor from torturing the child, a reluctance to interfere with the captor's autonomy should not be one of them. It seems to make *no* difference when we think about whether we ought to interfere in some very wrongful causal process (like torture) that doing so will involve interfering in the autonomy of an agent who is acting wrongly.

Let's look at what happens even if we grant Tooley the importance of this principle. Assume, as Tooley does, that all cases of interfering with the initiation of a causal process will involve interfering with someone else's actions. After all, the

⁹ Michael Tooley 'Abortion and Infanticide', *Philosophy and Public Affairs* 2 No. 1 (Autumn 1972) pp. 37–65, p. 58

point is that someone else is *causing* the process, which suggests that someone else must be doing something. Tooley suggests a distinction between a case where one has to interfere with someone else's actions, and a case where one does not. But what happens when we consider cases where not only the *causing* but also the *sustaining* of a process involves someone else's actions? Can Tooley make a further distinction between self-sustaining processes like poisoning, and processes that essentially require the continued action of another agent?

Compare the following two causal processes, one of which requires sustained action on the behalf of the initiator, and one of which does not.

- Pram:*
- (1) Anna intentionally pushes a baby's pram, which then freewheels down a hill where it will land in a lake if not stopped. Betty watches, and does nothing.
 - (2) Anna intentionally pushes a baby's pram across a flat park. Anna will push the pram into a lake if not stopped. Betty watches, and does nothing.

In (1), no sustained action is required on Anna's behalf. In (2), however, the causal process essentially requires sustained action to bring about the same result. Surely it cannot be the case that it is morally better for Betty to stand idly by in (2) than in (1)? It seems to me that Betty's actions in both cases are on a moral par. Doing nothing in (1) is equal to doing nothing in (2).

Yet if Tooley were to add a clause about interfering with an agent's action, he would be committed to drawing a moral distinction between Betty's non-interference in (1) where the pram is freewheeling, and Betty's non-interference in (2) where the pram is actively pushed. In addition, in (1) Tooley will have to equate idle Betty, who does nothing, with Anna, who deliberately causes the pram to freewheel down the hill. But in (2) he cannot equate idle Betty with Anna who deliberately pushes the pram across the park. Betty's inaction in (2) will somehow have to be morally better than her inaction in (1).

Not only does this asymmetry seem unintuitive to me as an advocate of the killing / letting die distinction. It also seems like it would be unappealing to a critic of the distinction. Why would someone who sees no difference between letting a person die and killing a person want to draw such a fine-grained distinction as that between

inaction when a pram is deliberately freewheeled into a lake, and inaction when a pram is deliberately pushed into a lake?

Even if Tooley could offer some kind of rationale for this principle, it would be another step altogether to make it appealing. If “the mere fact that when one prevents someone else from doing something one is interfering with someone’s action... is morally relevant”, it should be relevant in cases like *Pram*. And if it isn’t, we need some explanation of why it isn’t.

(iv) Reconstructing the Moral Symmetry Principle

So where is all this going? To remind ourselves: Tooley’s conception of events looked like this:

(1) Initiating the causal process

was equivalent to:

(2) Failing to interfere with the existing causal process.

But (1) was *not* equivalent to:

(3) Failing to prevent someone else initiating the causal process

But Tooley’s attempt to distinguish (2) and (3) has proved unsuccessful. It seems that (2) is equivalent to (3). But of course, if (2) is equivalent to (3), and (2) is equivalent to (1), (3) must also be equivalent to (1). So,

(1) Initiating the causal process

is equivalent to

(3) Failing to prevent someone else initiating the causal process.

What follows from equating a failure to prevent someone else from, for example, poisoning the whiskey with poisoning the whiskey oneself? Why does Tooley want to resist this identity? Well, if (1) and (3) are equivalent, then the inverse should similarly correspond.

(1) Not initiating the process of poisoning the whiskey*

is equivalent to

(2) Preventing someone else from initiating the process of poisoning the whiskey.*

If there is a moral obligation not to poison people, and this obligation is morally equivalent to an obligation to prevent *other people* from poisoning people, it looks very much as if I have just as stringent a duty to stop other people committing murder as I have not to commit murder myself. This is not a position that Tooley wants to endorse. Being just as responsible for the things that other people do as we are for the things we do ourselves is not an attractive view, and it certainly not a view that Tooley has argued for.¹⁰ But unless Tooley can plausibly restrict applications of the Moral Symmetry Principle to parts of causal processes that do not involve the actions of others, it seems that will be an implication of his argument.

III. Tooley and external variables

Tooley also attempts to undermine the killing / letting die distinction in another way. He suggests that it is external factors that make cases of killing *appear* to be worse than cases of letting die. He believes that those who defend the killing / letting die distinction have confused the ‘badness’ of external factors, like motivation, with the badness of killing and letting die themselves. Normally the circumstances surrounding instances of killing are worse than those surrounding instances of letting die. This, Tooley claims, is mistakenly taken to show that killing is *itself* therefore worse than letting die. I will argue here that Tooley’s account of the insignificance of the killing / letting die distinction legitimises the killing of innocent bystanders as a means to saving one’s own life.

(i) Tooley’s Machine Case

Tooley gives the following example to show how, when all external variables are removed, one cannot cite any moral difference between killing and letting die in themselves. We will call this *Machine*.

Machine: Two children, John and Mary, are trapped inside a machine, each in separate chambers. Above Mary is a canister of lethal

¹⁰ Tooley does acknowledge that if the Moral Symmetry Principle could not distinguish between (a) initiating a process oneself, and (b) failing to prevent someone else initiating that process, then this kind of agent-neutrality would indeed be implied by his view (p. 106). But since he takes himself to have shown a distinction between (a) and (b), he fails to take this objection seriously. I have undermined Tooley’s grounds for distinguishing between (a) and (b). If I am right, the Moral Symmetry Principle is subject to the standard objection to denying the killing / letting die distinction.

gas, which will soon be released into her chamber. However, by pushing a button, Passer-By can alter the machine so that the canister switches to John's chamber.

Tooley describes the position as follows.

If one pushes a button, John will be killed, but Mary will emerge unharmed. If one does not push the button, John will emerge unharmed, but Mary will be killed. In the first case one kills John, while in the second case one merely lets Mary die. Does one really wish to say that the action of intentionally refraining from pushing the button is morally preferable to pushing it, even though exactly one person perishes in either case? The best action, it seems to me, would be to flip a coin to decide which action to perform, thus giving each person an equal chance of surviving. But if that isn't possible, it seems to me a matter of indifference whether one pushes the button or not.¹¹

His claims about *Machine* suggest that Tooley thinks it permissible to redirect a harm onto a bystander in order to save a person who will otherwise be similarly harmed. But the fact that John has physical proximity to Mary doesn't seem like a good enough reason to allow Passer-By to kill him to save Mary's life. By putting John in the machine, Tooley implies that John is somehow involved in the lethal sequence of events in a way that does not make it unfair to include him in our moral equation. Mary's interests can thus be permissibly weighed against John's interests. According to Tooley, we can permissibly reduce the risk to Mary by transferring some of that risk to John when we toss the coin. But John is *not* part of the lethal sequence of events. The threat is to Mary's life, and to her life alone. John's physical proximity to Mary does not somehow make him a legitimate alternative target, whom we can use to reduce the threat to Mary's life.

(ii) *Harm, and risks of harm*

To see the problem with Tooley's argument, we need to consider the distinction between being able to avoid a harm, and being able to avoid being *at risk* of that harm. There is a difference between randomising a harm between John and Mary

¹¹ Michael Tooley 'An Irrelevant Consideration: Killing Versus Letting Die', p. 108

when *both* Mary *and* John are at risk of harm, and randomising a harm when only Mary is at risk.¹² Compare *Disease* with *Diseased Mary*:

Disease: John and Mary both have a fatal disease. There is only one dose of antidote. Neither John nor Mary have any prior claim upon the antidote. A third party must decide who gets the antidote.

Diseased Mary: Mary has a fatal disease. She can save herself only by infecting John (say by transmitting an errant parasite to him). John will then die, and Mary will survive.

In *Disease*, one ought to toss a coin to see who gets the last dose of antidote. In this case, randomising seems appropriate. But in *Diseased Mary*, one should not toss a coin to see whether Mary can permissibly infect John. This is true despite the fact that tossing the coin in *Diseased Mary* would give the children an equal chance of survival.

I think that we can explain the difference between *Disease* and *Diseased Mary* by thinking about the courses of action that we could justify to John. In *Disease* it is reasonable to think that both John and Mary would agree to a third party's tossing the coin, because they both have an interest in getting the antidote. It also seems reasonable that they would prefer that the choice be made fairly, if they have no indication of how the third party would otherwise choose. At least tossing the coin gives each a fifty-fifty chance of getting the antidote, whereas allowing the third party simply to choose may not.

But in *Diseased Mary* John has no reason to agree to a third party's tossing the coin to decide whether Mary can infect him with the parasite. He can gain no benefit from the randomisation, but might suffer a very grave harm. We might sometimes require a person to suffer a lesser harm in order that some innocent person be saved from a greater harm. But in *Diseased Mary*, the harm that will be suffered by John should he lose the toss is equal to the harm that is currently directed at Mary. It seems

¹² I take it to be the case that only Mary is at risk because, as things stand, the probability that Mary will be killed is 1 and the probability that John will be killed is 0.

to me that it would be wrong in *Diseased Mary* to force John to take on a risk of great harm when he can derive no benefit from that risk.

There are some cases in which it *would* be permissible for one to defend oneself against a threat to one's life by redirecting harm onto a bystander. Consider *Body Armour*:

Body Armour: Aggressor is shooting at Victim. Bystander is in the vicinity of Victim, and can protect herself by putting on protective clothing. Bystander refuses to put on the clothing, because protective clothing is unflattering. Victim can save himself only by deflecting a bullet towards Bystander.

I think it permissible that Victim intentionally deflect the bullet towards Bystander if this is the only way Victim can save his own life. In cases like *Body Armour*, Bystander not only has a chance to avoid harm, but a fair 'prior' chance to avoid even being *at risk* of harm. If Bystander declines to take that chance, it is permissible for Victim to redirect harm towards Bystander in defence of his life.¹³

There are questions about what counts as a 'fair chance' to avoid being at risk of harm. I am not going to attempt here to demarcate exactly what counts as fair, but rather suggest that there are some chances which are clearly fair, and some which are clearly not. I admit that there will be plenty of cases where it is hard to establish whether a person has had a fair chance: the scope of the grey area will be substantial. But there will also be cases on either side of the grey area where it is *not* difficult to see that a person either had, or did not have, a fair chance to avoid harm. If a person is informed of a likely risk and can take steps to avoid this risk at little cost to herself, this seems to count as her having a fair chance to avoid being at risk.

I am not suggesting that given one can avoid traffic accidents if one stays at home, all voluntary pedestrians can be viewed as having failed to avail themselves of fair chances to avoid being at risk of traffic accidents. The cost of never leaving the

¹³ I count *Body Armour* as an example in which the bystander had a chance of avoiding being at risk of harm, since in putting on the armour the bystander would no longer have been at risk. This seems different to *Diseased Mary*. If a third party decides to toss a coin, John finds himself at risk of being harmed without having had a fair chance to avoid that risk. I argue that we cannot kill bystanders who had only the chance to avoid the harm – they must also have had the chance to avoid the risk.

house is sufficiently great that we cannot regard those who decline this option as having declined a *fair* chance to avoid being at risk of harm. We might also think that while traffic accidents are a risk, they are not, in many circumstances, a *likely* risk. But it is sufficient to support my point that there are *some* occasions when people do have a fair chance to avoid harm and know the possible implications of not doing so, and their failure to do so has implications for how others may treat them.

This difference in treatment can be justified by the fact that by refusing to avoid placing themselves at risk of harm, a bystander can place additional burdens on other agents. The presence of a bystander can reduce the morally permissible options that are open to a person who is under threat of some harm. If the bystander had no chance to avoid her position, then it is impermissible to kill the bystander in the course of self-defence.¹⁴

But it seems wrong that a ‘willing’ bystander, i.e., one who could have avoided her position, should be afforded the same immunity as a bystander who had no chance to avoid her position. We might still think that if in *Body Armour* one could deflect the bullet in a different direction, away from the willing bystander, then one ought to do so. The willing bystander still merits this much consideration. But if one can make the alternative deflection only at much greater risk to oneself, I do not think that one is required to do this. One should not be required to bear a significantly

¹⁴ One of the anonymous referees pointed out to me that anyone who thinks that a bystander can flip the switch in a Trolley Problem, such that the trolley kills one instead of two, will think that the bystander can do this even if the one has no chance to avoid being at risk of harm. I do not think that the bystander is permitted to flip the switch in this case. There is insufficient scope here to expound my view on the Trolley Problem. But, briefly, those who think that flipping the switch is permitted might think about whether bystander ought to flip the switch in a case of one versus one. If he should not, this suggests that what is really going on in the two versus one case is a principle about numbers *overriding* a principle about killing and letting die. It does not show that the killing / letting die distinction is irrelevant. We might also think that there is a difference between what the two towards whom the trolley is currently heading can do to the bystander in self-defence, and what the one towards whom the trolley would be redirected can do to the bystander in self-defence. It seems to me that it would be impermissible for the two to shoot the bystander to make his body fall on the switch, altering the course of the trolley towards the one. But it would *not* be impermissible for the one to shoot the bystander in self-defence if the bystander moves to flip the switch to alter the trolley’s course so that it will kill the one. See Section III.(iii) below. Importantly, our intuitions in these cases seem to be radically altered if we add that the one saw the trolley approaching the five but wandered onto the alternative track anyway because it was quicker than using the bridge. Here, the one had a fair chance to avoid being at risk of harm, and has placed greater burdens on those whom the trolley approaches. That it seems much clearer *here* that we may redirect the trolley away from the two towards the one supports my claim that there is a difference between chances to avoid harm and chances to avoid being at *risk* of harm.

greater burden than one would have had to bear had the bystander taken her fair chance of avoiding being at risk of harm.¹⁵

So, it seems to me that tossing the coin in *Machine* would be like tossing a coin to determine whether Mary can permissibly infect John in *Diseased Mary*. In both cases, John has no fair chance to avoid being at risk of harm, and no reason to agree to the randomising of the harm. Tooley's suggestion of tossing a coin to even the odds is out of place here. One can permissibly kill only those bystanders who had a fair chance to avoid being *at risk* of harm. One cannot permissibly kill bystanders who have only a fair chance of avoiding the harm itself.

Even if pushing a button that will kill John is the only way that Mary can save her own life, it is impermissible for her to do so. And it cannot be permissible for Passer-By to kill John on Mary's behalf if it is impermissible for Mary to kill John herself. If *she* cannot push a button to redirect a threat to her life towards a bystander, it cannot be that case that someone else can push it for her.

This, I think, casts some serious doubt on Tooley's claim that it is a matter of moral indifference whether or not I push the button. Redirecting a lethal threat towards an innocent bystander is never a matter of moral indifference. But it may not be obvious that a prohibition on killing bystanders supports a distinction between killing and letting die. We could have other reasons for believing in the wrongness of killing bystanders that do not rest upon the killing / letting die distinction.

(iii) Self-defence and responsibility

One of the most effective ways that we can show that the killing / letting die distinction is at work in *Machine* is by thinking about what John and Mary can permissibly do in self-defence. Imagine that Passer-By decides to push the button that will kill John. However, John, being a forward-thinking person, has a shotgun concealed about his person. Can John permissibly shoot Passer-By, who poses a lethal threat to him, in order to stop Passer-By from pressing the button? It seems to me that he can.

Now consider whether it is permissible for Mary to use harmful means to force Passer-By to push the button. If, on the one hand, we answer that she can, in

¹⁵ There might be a case for claiming that a willing bystander as described in *Body Armour* is not really a bystander at all, on the grounds that she culpably reduced the potential victim's morally permissible options. But this adjustment will not affect my thesis – it will still support the claim that we can only kill those who had a chance to avoid being at risk of harm (whether we call them bystanders or not).

order to preserve the similarity between the two cases, we are committed to the implausible claim that one may use seriously harmful means to force a person to come to one's aid at the cost of someone else's life. If, on the other hand, we answer that Mary may *not* use such force, this supports a significant difference between killing and letting die. John may use potentially lethal force to defend himself against a person who is going to kill him, *even though* someone else (Mary) will die as a result. But Mary cannot use such force against a person who intends to let her die, *especially* if someone else will die if she succeeds.¹⁶

We can further support the relevance of the killing / letting die distinction by envisaging *Revised Machine*.

Revised Machine: John and Mary are again in separate chambers, but now Passer-By cannot see the gas canister. Thus, Passer-By does not know at which chamber the canister is currently pointed. Passer-By knows only that pushing the button will cause the canister to *switch* chambers.

So, Passer-By has no idea whether pushing the button will kill John or kill Mary. The proviso against tossing a coin still holds in the revised case, and I think that intuitions against pushing the button are perhaps clearer here than in the original *Machine* case. Given that Passer-By does not know which child will die either way, it seems better that he do nothing. For what he *does* know is that in pushing the button he will kill a child, whereas in not pushing the button he will let a child die. It seems implausible that one would opt for *causing* a child's death over *allowing* a child's death if one does not know which child will be killed. This is true despite the fact that in terms of the number of deaths that result, it makes no difference whether or not Passer-By pushes the button, since one child will die either way.

Where it does make a difference whether or not Passer-By presses the button is in terms of the responsibility that he will bear for a child's death. We can hold constant

¹⁶ Frances Kamm advances a similar line of argument against Warren Quinn in *Morality Mortality Volume II: Rights, Duties and Status*, OUP: New York (1996) esp. pp. 76 – 79, 101 – 104. Bear in mind that it does not have to be wrong in general for a person to threaten to harm someone who refuses to save them. It may even be permissible to exact some harm upon someone who refuses to save you. But that a person may not do *as much* against someone who refuses to save them as they may do against someone who is going to kill them is sufficient to underline the killing / letting die distinction.

all the variables that Tooley cites important, such as motivation and outcome, and still point to a difference between the case where Passer-By kills John, and the case where he lets Mary die. If he kills John, Passer-By breaches a more serious responsibility that he owes to John than he breaches if he lets Mary die.

That killing John breaches a more serious responsibility than letting Mary die seems to stem from the fact that John would have survived had Passer-By been absent. Whether some event would have occurred whether or not I had been present seems, at least *prima facie*, to bear on my responsibility for that event.¹⁷ In Frances Kamm's words,

In dying, the [victim of letting die] loses only life he would have had via the agent's help. By contrast, the standard [victim of killing]... loses life he would have had independently of the killer.¹⁸

I think that we can perhaps be more specific about the difference that Kamm highlights. If Passer-By lets Mary die, Mary is no worse off for Passer-By's presence than she would have been for his absence. Indeed, Mary has a possible lifeline only *because* of Passer-By's presence. Thus Passer-By cannot be said to make Mary worse off for being there and failing to save her. But if Passer-By were to kill John, John is undeniably worse off for Passer-By's presence.

Conclusion

I have argued that Tooley's rejection of the killing / letting die distinction is ultimately unsuccessful. I demonstrated that the Moral Symmetry Principle, as defended by Tooley, is open to the standard criticism that a denial of the killing / letting distinction will equate acts that are intuitively different. Tooley's suggested principle of interfering in other people's actions was shown to fail. Thus, the Moral Symmetry Principle cannot distinguish between failing to prevent other people from causing an event and causing that event oneself. The implication of this equation is that I am just as responsible for preventing others from causing harm as I am for not causing harm myself.

¹⁷ I say *prima facie* because we could think of cases in which, say, you promise to shoot Victim unless I shoot Victim. That Victim would have died whether or not I had shot him does not seem to alter the fact that I am responsible for his death if I do shoot him. But these tricky cases of over-determination are, I think, insufficient to undermine the general principle.

¹⁸ Frances Kamm *Morality, Mortality Volume II: Rights, Duties, and Status* p. 23

I rejected Tooley's contention that there is no moral difference between letting Mary die and killing John in *Machine*. I argued that we must distinguish between bystanders who have a fair chance to avoid harm, and bystanders who have a fair chance to avoid being *at risk* of harm. Only those bystanders who have a chance to avoid being at risk of harm are liable to be killed in self-preservation. Tossing a coin in *Machine* places John at risk of a harm that he had no opportunity to avoid, and from which he can derive no benefit.

That John can do more in self-defence against Passer-By than Mary can do against Passer-By further underlines the relevance of the killing / letting die distinction. John is permitted to kill a person who is going to kill him. Mary cannot permissibly kill a person who refuses to save her, especially if the only way she can be saved involves the killing of John. That John is made worse off by Passer-By if Passer-By presses the button is also significant. Passer-By does not make Mary worse off by his presence if he fails to save her, since it is only *because* Passer-By is present that Mary can be saved at all.

References:

- Bennett, J - *The Act Itself*, OUP: New York (1995)
- Kagan, S - *The Limits of Morality*, OUP: Oxford (1989)
- Kamm, F M - *Morality, Mortality Volume II: Rights, Duties, and Status*, OUP: New York (1996)
- Rachels, J - 'Active and Passive Euthanasia', *Killing and Letting Die*, eds. Alastair Norcross & Bonnie Steinbock, Fordham University Press: New York (1994)
- Tooley M - 'Abortion and Infanticide', *Philosophy and Public Affairs*, Vol. 2 No.1 (Autumn 1972) pp. 37 - 65
- 'An Irrelevant Consideration: Killing Versus Letting Die', *Killing and Letting Die*, eds. Alastair Norcross & Bonnie Steinbock, Fordham University Press: New York (1994)