

Original citation:

Tadros, Victor. (2017) Permissibility in a world of wrongdoing. *Philosophy and Public Affairs*, 44 (2). pp. 101-132.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/84022>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"This is the peer reviewed version of the following article Tadros, Victor. (2017) Permissibility in a world of wrongdoing. *Philosophy and Public Affairs*, 44 (2). pp. 101-132., which has been published in final form at <http://doi.org/10.1111/papa.12074> . This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#)."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Permissibility in a World of Wrongdoing*

VICTOR TADROS

School of Law, Warwick University

Sometimes X's conduct (by which I mean both her acts and omissions) may result in bad consequences because of Y's wrongdoing. Does this affect the permissibility of X's conduct?

One way in which it might is that by acting wrongly Y may forfeit certain rights. This affects the duties X owes to Y. Similarly, if X and Y owe duties of reciprocity to each other and Y wrongly fails to do her duty, X need not do what she would otherwise have been required to do. I leave these questions aside.

Our question is whether Y's wrongdoing makes a difference to what X is permitted to do because of what X owes to others. I focus especially on X's duties to protect others from harm, and not to harm others. And I focus only whether Y's wrongdoing makes a difference to the permissibility of X's conduct as such. Y's wrongdoing makes a difference to these things as such

* I am to audiences at Oxford, Rutgers and University College London. Thanks also to Johann Frick and his graduate class at Princeton, Tom Hurka, George Letsas, Jeff McMahan, Mike Martin and two Associate Editors at *Philosophy and Public Affairs*. I am also the grateful recipient of a Major Research Fellowship from the Leverhulme Trust, which afforded me the time to work on this paper.

only if it makes a difference when we hold other things equal, including the magnitude of harm that people will suffer and who will suffer it. Thus I leave aside, for example, the effects of the incentives X would create by not responding to Y's wrongdoing. I also assume that Y cannot be prevented from acting wrongly, and cannot be made to pay compensation either to X or to others for her wrongdoing.

Take a pair of cases where X can either prevent V from being harmed, or cause V harm. In the first case, the harm that V suffers will result from Y's wrongdoing; in the second it will not. The harm that V will suffer is held equal, as is everyone's future conduct. Y cannot be prevented from acting wrongly, nor can she be made to compensate either X or V. The idea that Y's wrongdoing makes a difference as such implies that whether it is permissible for X not to prevent harm to V, or to harm V, may differ between these cases.

I show that, other things equal, the fact that the harm that results from X's conduct is attributable to Y's wrongdoing may make a difference to the permissibility of X's conduct. But if it does the difference it makes is typically not the one that many defend: that X is permitted not to prevent V from being harmed, or to harm V, because the harm that V will suffer will be Y's responsibility. On the contrary, although the expected conclusion may be defensible in certain circumstances, the fact that V will be harmed as a result of Y's wrongdoing often requires X to prevent V from being harmed, or to refrain from harming V, where otherwise she would lack such a duty.

I show this by exploring the effects of Y's wrongdoing on the stringency of X's duty to prevent V from being harmed, or not to harm V. Here is why the stringency of our duties determines what we are permitted to do, all things considered. X's *prima facie* or *pro tanto* duties, such as the duty to prevent harm, or the duty not to harm others, can be vitiated or overridden in particular circumstances where X can prevent the harm, or not cause it. They are vitiated or overridden if the conduct described in the *prima facie* or *pro tanto* duty would, in those circumstances, be too costly either to X or to others. The less stringent X's *prima facie* or *pro tanto* duties, the more easily these duties are vitiated or overridden. When X's duties are vitiated or overridden she is permitted to do what her *prima facie* or *pro tanto* duty would otherwise have required.

For example, if the duty to prevent harm to others is very stringent, X is all things considered required to prevent harm to another even when the costs that she will bear, or that others will bear, if the rescue is performed is high. If it is less stringent, X is all things considered permitted not to prevent harm to another if the costs that she will bear, or that others will bear, if the rescue is performed are high. I suggest that X's duty to prevent V from being harmed, or not to harm V, is typically either unaffected, or is strengthened, by the fact that the harm that V will suffer will result from Y's wrongdoing. I will also show why exceptions to this general idea arise.

This issue has great practical significance. For example, in 2015 the UK cut development aid to India, despite the fact that many Indians live in dire

poverty. Some in the UK objected to providing aid because the Indian government spends a great deal of money that it could spend helping its poor on its space programme. One argument in favour of cutting aid is that if the UK does so, the resulting poverty will result from the wrongdoing of the Indian government, and will be their responsibility not the UK's.

More generally, if the burdens of poverty relief were widely shared amongst the wealthy, poverty might be alleviated at little cost to anyone. But many wealthy people do not pay their fair share. Are the rest of us required only to pay our fair share, or must we pay more to make up the shortfall left by those who will not pay?

Or consider innocent people who are put in harm's way by wrongdoers. Israel's targeted attacks on Hamas militants in Gaza have seemed to many disproportionate because of the number of innocent people killed as a side-effect. Israeli politicians and supporters respond that those they were targeting were terrorists who wrongly hid amongst civilian populations, effectively using them as human shields. This, they argue, makes their actions proportionate.

Or consider wrongful intervening agency in the US led war in Iraq that commenced in 2003. Roughly, wrongful intervening agents are those people who are independent of US forces, and who acted wrongly as a result of the war, causing harm. Many people think that the number of innocent civilians killed made the war disproportionate. Some innocent civilians were killed directly by US combatants and their allies. Insurgents who wrongly

exploited the chaos of the war to make inroads into Iraq killed others. Many people believe that the deaths caused by insurgents count for less in the proportionality calculation than those caused directly because the insurgents are wrongful intervening agents.

Even those who are sceptical about these cases might find this one compelling: Gandhi's non-violent struggle for Indian independence predictably resulted in violence, especially due to the brutal police response. During the famous Salt March, several people were beaten, wounded and killed. Gandhi was reported to have said:

My heart now is as hard as stone. I am...ready to sacrifice thousands and hundreds of thousands of men if necessary. In this game of dice we are playing the throw has been as we wanted. Should we then weep or smile?¹

This sentiment might seem less appealing if those who were injured or killed were trampled underfoot by Gandhi and the other marchers, even if this would have been equally effective in securing independence. But Gandhi might have thought that as the police were responsible for the beatings and the killings, they did not turn his protest into a violent one in the morally relevant sense.

¹ S Wolpert *Gandhi's Passion: The Life and Legacy of Mahatma Gandhi* (Oxford: OUP, 2001) 148.

I proceed as follows. Sections I, II, and III are concerned with the duty to prevent harm to others. I reject the two most plausible arguments to support the view that non-compliance makes this duty less stringent, and suggest that there is some support for the opposite conclusion. Sections IV and V are concerned with acts that harm others. Section IV considers whether it makes a difference that those who will be harmed by our actions have wrongly been put in harm's way, and concludes either that it does not, or that our duties not to harm others are more stringent in that case. Our duties not to harm others may be less stringent if others have been put in harm's way with the intention of harming the duty holder or others, or to facilitate a further wrong that the person intends by making it wrong for us to prevent that further wrongdoing. But even these ideas have only modest moral significance. Section V considers whether cases of intervening agency should be treated differently from cases where others have wrongly been put in harm's way, and concludes that they should not.

I. Duties of Assistance and Fair Shares

Assume a moderate view of the duty to prevent harm to others where no one has acted wrongly: we have a duty to bear non-trivial costs to prevent harm to others, but we need not do so in all cases where doing so would achieve the best outcome considered impartially. To make our discussion more precise, consider the duty to prevent a person from being killed where no one else's

wrongdoing will be implicated in the person's death if she is not rescued: other things equal, X has a duty to prevent V from being killed if he can do so at some cost less than n , but not if he can do so only at a cost greater than n , where n is non-trivial, but is not equal to or greater than the harm that X will prevent by saving V. n is the maximum cost that X must bear to prevent V from being killed.

Now suppose that a group of people, including X, can together prevent V from being killed by sharing a cost between them where the group is required to share the cost. However, the other members fail to do their fair share. What is X required to do? It will help to compare the duty that X would have had if the non-compliers had been unable to help. If they were unable to help, they would not act wrongly by failing to do so (assuming some version of 'ought implies can'). Our question, then, is whether non-compliance makes a difference when compared with inability, other things equal.

One unpopular view is that X's duty is more stringent in cases of non-compliance than cases of inability. A more familiar view is that non-compliance makes no difference:

Take up the Slack: other things equal, if X ought to share the cost of preventing harm to V with others, and these others fail to do their fair share, X's duty to assist V is as stringent as it would have been had these others been unable to assist.

Alternatively, some argue that X's duty is set by the costs that he would have been required to bear were the non-compliers to have done their fair share:

Fair Share I: other things equal, if X ought to share the cost of preventing harm to V with others, and these others fail to do their fair share, X's duty is only as stringent as it would have been had these others done their fair share.²

Fair Share I is the view that X lacks a duty to bear a cost that is greater than the cost that he would have been required to bear had these others done their fair share. This implies that X need not prevent V from being harmed even if the cost that he would bear is significantly less than the cost that he would have been required to bear had the non-compliant members of the group been unable to contribute.

We can dismiss this view without further ado by considering:

² Roughly this is defended in Murphy *Moral Demands in Nonideal Theory* (Oxford: OUP, 2000) ch.5. In chapter 7, Murphy considers whether rescue cases are special. He is inclined to think that they are not. David Miller, in contrast, thinks that we have humanitarian obligations to take up the slack (though he is unclear about how stringent these obligations are), but no duties of justice to do so. The difference between these duties is, for him, primarily a matter of enforceability. See 'Taking Up the Slack: Responsibility and Justice in Situations of Partial Compliance' in C Knight and Z Stemplowska *Responsibility and Distributive Justice* (Oxford: OUP, 2011).

Boat: V is drowning in a pond. He can be rescued by paying £20 to release a boat from an automatic locker. X and Y each have plenty of money, but as she does not care for V, Y refuses to contribute £10.

X and Y ought to share the cost of rescuing V between them and pay £10 each. Suppose that had Y been unable to contribute, X would have been required to rescue V at the cost of £5000. X is surely required to pay £20 to save V's life. Even if Y's failure makes some difference to the stringency of X's duty, it cannot make that much difference.³

A more plausible view is that non-compliance makes a difference to the stringency of the duty to prevent harm:

Fair Share II: other things equal, if X ought to share the cost of preventing harm to V with others, and these others fail to do their fair share, X's duty to prevent harm to V is less stringent than it would have been had these others been unable to contribute.

This need not implausibly imply that X is permitted not to save V's life in *Boat*. X's duty is less stringent than it would have been had Y been unable to contribute, but it may be sufficiently stringent to require X to rescue V at an

³ For somewhat similar objections, see also G Cullity *The Moral Demands of Affluence* (Oxford: OUP, 2004) ch.5.2 and Z Stemplowska 'Doing More than One's Fair Share' (2016) 19 *Critical Review of International Social and Political Philosophy* 591.

additional cost of £10. *Fair Share II* has implications for permissibility in cases where X is closer to the maximum cost that he would have to bear to meet a need.

Compare:

Pond I: Through no one's wrongdoing V is drowning in a pond. X can rescue V at cost n . If X does not do this, V will die. Y cannot assist.

Pond II: V is drowning in a pond. X and Y can together rescue V at a very small cost to each, fs . X is willing to do his fair share. Y refuses to contribute, and cannot be forced to do so. X alone can rescue V at cost n (which is $>fs$). If X does not do this, V will die.

V will die if X does not bear a cost n in both cases. The difference between them is that Y's wrongdoing makes it necessary for X to bear cost n , rather than the lesser fs , to prevent the harm in *Pond II*, but not *Pond I*.

Fair Share II implies that X's duty to rescue in *Pond II* is less stringent than his duty to rescue in *Pond I*. Now recall that n is the maximum cost that X would be required to bear to rescue V in *Pond I*. As Y's non-compliance in *Pond II* weakens X's duty when compared with *Pond I*, *Fair Share II* thus implies that X is not required to rescue V in *Pond II*.

II. Responsibility

How can *Fair Share II* be defended? Here is one rough idea that some might find attractive. If X does not rescue V in *Pond II*, Y will bear primary responsibility for V's death. This is because Y could ensure not only that V is saved, but also that V is saved with no one acting wrongly. Because Y will be primarily responsible for V's death, X will bear less responsibility for V's death if he does not rescue V than he does in *Pond I*.

The Responsibility Argument is that as X will bear less responsibility for V's death if he does not rescue V in *Pond II* than in *Pond I*, his duty to rescue V is less stringent in *Pond II* than in *Pond I*.

There are different senses of responsibility. One is forward-looking. For example, parents are responsible for looking after their children; police officers have responsibility for law and order; I am responsible for controlling my own conduct; and so on. This sense of responsibility cannot feature in an argument for any *Fair Share* principle, though, because claims about what we are responsible for doing in this sense are too close to claims about our duties. The view that X is not responsible for rescuing V in *Pond II* because this is Y's responsibility is more or less the view that X does not have a duty to rescue V because Y has the duty to do so. And that more or less just restates the claims of friends of *Fair Share* rather than defending them.

Another sense of responsibility is backward-looking: roughly, my responsibility for some act or outcome depends on how deeply that act or

outcome is attributable to me as an agent, or as a person. *The Responsibility Argument* most plausibly rests on this sense of responsibility.

Here is a general statement of the argument thus understood:

If:

- i) X and Y have a duty to share the costs of *a*ing; and
- ii) They are required fairly to share the burden of doing *a* between them; and
- iii) X is willing to comply with this duty; but
- iv) Y will not comply; and
- v) X can do *a* alone, but at greater cost than he would have borne had the burden been shared with Y.

Then

- 1) If *a* is not done, Y will bear greater responsibility than X for *a* not being done.
- 2) Y's greater responsibility for the *a*ing not being done lessens X's responsibility for it.
- 3) Therefore X's responsibility for *a*ing not being done is less than it would be were Y to have been unable to contribute to *a*, and thus was not responsible for *a*ing not being done.

- 4) The stringency of X's duty to *a* depends on how responsible he will be for *a*ing not being done.
- 5) Therefore X's duty to *a* is less stringent than it would have been had Y been unable to contribute to *a*.

There are two objections to this argument. First, it does not seem true that one person's responsibility for an act or outcome diminishes another's. Compare these two cases:

Demolition 1: X works for a demolition firm. He discovers that V has entered a house that he is about to demolish, but as he is on a tight schedule he demolishes the house, killing V.

Demolition 2: As *Demolition 1*, except Y has wrongly forced V into the house in order that X will kill V.

Y is fully responsible for V's death in *Demolition 2*. Furthermore, only Y could ensure that everyone complies with his duties. X has a contractual obligation to demolish the house. Y has a duty not to force V into the house. Had Y complied with his duty, X would have complied with his duty as well. X does not intend Y to kill V; he would prefer that Y did not do this. Yet X's responsibility is not plausibly diminished in *Demolition 2* when compared with *Demolition 1*.

Now suppose that this is wrong, and that one person's responsibility for an outcome diminishes the responsibility of others in the relevant cases. The stringency of a person's duty to rescue nevertheless seems unaffected by the responsibility of others for the outcome of his failure to rescue. To see this, consider:

Pond III: Y has wrongly knocked V into a pond. Y can rescue V but will not do so. X and Y can together rescue V at some small cost to each. Y will not help. X can rescue V at cost n . If he does not do this, V will die.

If X does not save V, Y is even more deeply responsible for V's death than he is in *Pond II*, plausibly assuming that one is more deeply responsible for causing harm than for failing to prevent it.

And the same conclusion is true even if we consider only the backward-looking responsibility that a person has for an outcome because of her failure to fulfill her positive obligations. Y should bear a much greater share of the costs of rescuing V than X, given that Y wrongly put V in peril. So if V is killed, the responsibility that Y has for the death due to her failure to fulfill her positive obligations is much greater than X's.

The Responsibility Argument thus implies that X's duty to rescue is less stringent in *Pond III* than *Pond I*. This is not a plausible view. Some people believe that we have no more reason to rescue people from wrongdoing than from natural disaster; others that our duties to rescue people from harmful

wrongdoing are more stringent than our duties to rescue them from non-wrongfully inflicted harm.⁴ But it is powerfully counterintuitive that our duty to protect people from wrongfully inflicted harm is less stringent than our duty to protect people from natural disaster as *The Responsibility Argument* implies.

III. Victims of Unfairness

Some defend *Fair Share II* by pointing to the fact that if we are required to take up the slack of non-compliers, we will be victims of unfairness. As we will see, when it is carefully stated this argument need not have the counterintuitive implication of *The Responsibility Argument* in *Pond III*.

Recall:

Pond I: Through no one's fault, V is drowning in a pond. X can rescue V at cost n . If X does not do this, V will die. Y cannot assist.

⁴ For discussion, see J McMahan 'Humanitarian Intervention, Consent, and Proportionality' and P Singer 'Bystanders to Poverty' both in N A Davis, R Keshen, and J McMahan *Ethics and Humanity: Themes from the Philosophy of Jonathan Glover* (Oxford: OUP, 2010); V Tadros *The Ends of Harm: The Moral Foundations of Criminal Law* (Oxford: OUP, 2011) and *Wrongs and Crimes* (Oxford: OUP, forthcoming) ch.9.

Pond II: Through no one's fault, V is drowning in a pond. X and Y can together rescue V at a very small cost to each, fs . X is willing to do his fair share. Y will not comply, and cannot be forced to do so. X alone can rescue V at cost n (which is $>fs$). If X does not do this, V will die.

Suppose X is required to rescue V in *Pond II*, that he complies with all his duties, and that he would not rescue V were he not so required. Y's non-compliance then results in X bearing a greater cost ($n - fs$) than he would have borne had Y done his fair share. He is not free, in the moral sense, to avoid this cost. X is then a victim of Y's wrongdoing. This might be thought to affect the stringency of X's duty to assist V. As X is a victim of Y's wrongdoing only if X has a duty to rescue V there is a reason against X having this duty.

Here is the argument more precisely:

The Infliction Argument:

If:

- i) X and Y have a duty to share the costs of a ing; and
- ii) They are required fairly to share the burden of doing a between them; and
- iii) X is willing to comply with this duty; but
- iv) Y will not comply; and

- v) X can do *a* alone, but at a greater cost than he would have borne had the burden been shared with Y.

Then:

- 1) If X is required to *a*, Y will have done something morally equivalent to inflicting a cost on X.
- 2) X will then be a victim of Y's wrongdoing in bearing the excess cost.
- 3) Had Y been unable to share the costs of *a*ing, and X is required to *a*, X is a victim of bad luck in bearing the equivalent of the excess cost.
- 4) There is a greater reason against X having a duty if having that duty will make him a victim of wrongdoing than if it will make him a victim of bad luck.
- 5) Therefore there is a reason against X having a duty to *a*.
- 6) Therefore, X's duty to *a* is less stringent than it would have been had Y been unable to share the cost of *a*ing.

This argument reflects the conviction that many people have that they are not required to take up the slack that non-compliers leave because the non-compliers will have wrongly imposed on them. For example, if some of colleagues fail to do their marking, many feel that others in the department

are not required to take up the slack because they would be victims of unfairness were they so required.

One doubt about this argument concerns premise 4). It is a reason against my having a duty to *v* that my *v*ing will be costly to me. But is the fact that I will be a victim of wrongdoing if I have a duty to *v* the kind of thing that can affect whether I have that duty? Answering this question requires a more general theory of wrongdoing that I cannot provide here. Premise 4) might seem more attractive on some theories than others. For example, it might seem especially attractive to those who believe that duties are the outcome of an idealized process of interpersonal reasoning aimed at norms of social regulation. Perhaps some who reject this view of morality might also be able to admit reasons such as those in premise 4) into evaluation of our duties. But some will argue that this is the wrong kind of reason to offer for or against a duty. I cannot resolve these deep disputes here.

Another objection to premise 4) is that it is no worse to be a victim of wrongdoing than to be a victim of bad luck. I noted earlier that some people think that there are no stronger reasons to prevent people from being victims of harmful wrongdoing than there are to prevent them from being victims of non-wrongfully inflicted harm. If *X* has a duty in both *Pond I* and *Pond II*, *X* is equally unlucky in having to bear a cost. The only difference is the source of the cost. There is no more reason to relieve *X* of this cost if he will be a victim of wrongdoing than if he will be a victim of bad luck. Therefore, there is no more reason against *X* having the duty to rescue *V* in *Pond II* than in *Pond I*.

This argument is also controversial. Not everyone agrees that there is equal reason to rescue people from harmful wrongdoing than from non-wrongfully inflicted harm. Assessing this issue requires us to address deep and controversial questions about the nature of wrongdoing, and our reasons not to want to be wronged.

Let us, therefore, grant premise 4). Even so, *The Infliction Argument* fails, for if premise 4) is true there is a countervailing moral difference between *Pond I* and *Pond II*. If X is not required to rescue V in *Pond II*, and in consequence does not do so, V is a victim of Y's wrongdoing. Had Y done what he was required to do, V would have been saved; as Y did not, V drowns. In contrast, if X is not required to rescue V in *Pond I*, and X does not do so, V is not a victim of Y's wrongdoing.

If we believe premise 4), we should also believe that there is a reason for X having a duty if his lacking that duty makes others a victim of wrongdoing.⁵ And not only that: the reason for X having this duty seems more powerful than the reason against. The wrong that Y would do to X were X required to rescue V results in X bearing a modest cost, whereas the wrong that Y would do to V were X not required to rescue V involves a much greater cost – for V, it is a matter of life or death.

And there is a further argument to support X having this duty. If X has this duty the overall cost that he is required to bear is no greater than the cost

⁵ For a somewhat related idea, see M Ridge 'Fairness and Non-Compliance' in B Feltham and J Cottingham *Partiality and Impartiality: Morality, Special Relationships, and the Wider World* (Oxford: OUP, 2010).

that he would have been required to bear had Y not been a wrongdoer because he was unable to share the cost, say because he was not there. Hence, if X is required to rescue V, X is not a victim of *there being a wrongdoer present* – had Y not been able to contribute, X would have been no worse off. In contrast if X is not required to rescue V, and consequently does not do so, V is a victim of *there being a wrongdoer present* – had Y not been able to contribute because he was not present, and X complies with his duties, X would have rescued V.

Here is a response. We must not assume, as this argument assumes, that all that matters is the magnitude of the burden that a person will suffer if she is a victim of injustice.⁶ We must take account of the way in which the burden is inflicted. If X is required to rescue V in *Pond II*, Y has done something morally equivalent to harming X. Furthermore, X is harmed through the exercise of his own agency – he is put in a position where he must harm himself. If X is not required to rescue V in *Pond II*, and does not do so, in contrast, Y has allowed V to die. Doing harm is worse than allowing

⁶ Ridge seems to make this assumption in calculating how the costs of non-compliance should be distributed in his *Reservoir Dogs* tipping case at 207-8. But he may intend to qualify that approach given his brief discussion of doing and allowing at 212-3.

harm. Hence, X will be a victim of a worse kind of wrongdoing if he is required to rescue V than V will be if X is not required to rescue V.⁷

This response has a further attraction. I earlier noted that *The Responsibility Argument for Fair Share II* implausibly implies that we have stronger reason to save victims of natural disaster than victims of wrongdoing. We can now see that *The Infliction Argument* need not have this implication. In *Pond III*, if X is not required to rescue V, and does not do so, Y has killed V rather than preventing V from being saved. Hence, *The Infliction Argument* has the plausible implication that X's duty in *Pond III* is more stringent than his duty in *Pond II*, and perhaps even more stringent than in *Pond I*.

Here is a reply to the response. Suppose that X is not required to rescue V and does not do so. X would have saved V if Y had done what he was required to do. Although Y does not harm V directly, he creates the conditions where X does not rescue V. This might be thought similar to preventing a person from being saved. Thus Y's conduct does not amount to a mere allowing.

However, preventing a person from being saved is also arguably less bad than harming a person. For example, it seems wrong to defend oneself against a culpable attacker if doing so will directly kill five people as a side-effect. But it seems permissible to defend oneself against a culpable doctor

⁷ This is how Ridge, at 212-3, interprets Murphy's discussion in *Moral Demands in Nonideal Theory* at 92. I am not sure that this is what Murphy intends.

even if this results in the doctor not saving five people. Some people conclude that there is an important moral difference between directly killing five and preventing five from being saved.⁸ For the same reason, it might be argued, if X is required to rescue V, X is a victim of a worse kind of wrongdoing than V will be if X is not required to rescue V.

But even understood in this way *The Infliction Argument* does not support *Fair Share II*. Even if it is worse to inflict harm than it is to prevent a person from being saved from harm *other things equal*, which is itself controversial, it is not worse to harm a person than it is to prevent a person from being saved when the harm is not equal. It is not worse to harm a person to a small degree than it is to prevent a person from being saved from a grave harm. If X is required to rescue V in *Pond II*, and he does what he has a duty to do, he suffers an extra cost, and Y inflicts that on him (or does the moral equivalent of inflicting it on him). If X does not rescue V, Y will have prevented V's life from being saved. And this seems much more gravely wrong.

The considerations that *The Infliction Argument* relies on, if anything, support the view that X has a more stringent duty to rescue V in *Pond II* than in *Pond I*. If he has such a duty and he acts on it, V's life is saved, and V is not victim of grave wrongdoing by Y – Y will not have created the conditions in which X does not save V. We have thus found no support for *Fair Share II* and

⁸ See M Hanser 'Killing, Letting Die and Preventing People from Being Saved' (1999) 11 *Utilitas* 277; J McMahan 'Proportionate Defense' (2013-14) 23 *Journal of Transnational Law and Policy* 1, 4-5.

some support for the unpopular view that we have more stringent duties to rescue others in cases of non-compliance than in cases of inability, at least when our duties are grounded in the need to prevent serious harm. It seems difficult to improve on *The Infliction Argument* to support *Fair Share II*. I think that the driving force behind the *Fair Share II* is the unfairness that X will suffer if he is required to take up the slack of Y's failure to do his duty. But once we clearly see that V will also be a victim of unfairness if X does not have this duty, and that the unfairness to V will be even greater, we see that the idea that X will be a victim of unfairness cannot support *Fair Share II* after all.

So far we have only considered cases of great and urgent need. Might *The Infliction Argument* have different implications in cases where the costs are reduced?⁹ Consider:

Advice I: X and Y have one research student each: V1 and V2. Each student is need of advice before a conference. Y refuses to give V2 the advice she needs. X can advise both V1 and V2, but this will be time consuming.

Advice II: As *Advice I* except that Y doesn't turn up to work because she is ill, and is thus unable to advise V2.

⁹ I am grateful to an associate editor of *Philosophy and Public Affairs* for encouraging me to explore this issue.

Suppose that X and Y have a duty equally to share the burdens of giving advice to research students. Furthermore, suppose that it is wrong for Y to fail to advise V2 in *Advice I*. And suppose that X is required to advise both V1 and V2 in *Advice II*. Finally suppose that all other things are equal – in particular that X's duty to advise both students in *Advice II* is not due to the fact that Y will reciprocate if X is ill, but simply because of the needs of the student. Might X lack a duty to advise V2 in *Advice I*?

In order to see how the premises of *The Infliction Argument* apply to this case we must consider the gravity of the wrong that X will suffer if she is required to advise V2, and the gravity of the wrong that V2 will suffer if X is not required to do so. This is difficult to do in this case. But it is at least commonly true that the burdens that our positive obligations impose on us are less than the benefits that others receive from our fulfilling them. This also seems true in this case: the benefit that supervisees receive from the advice of supervisors is greater than the burden that supervisors bear in giving it. Were this not so, giving the advice would normally be supererogatory. This suggests that the gravity of the wrong that V2 will suffer if X does not give the advice is greater than the gravity of the wrong that X will suffer if he is required to give the advice.

Against this, it might be argued that the fact that X needs to perform some positive act to benefit V2 looms larger in cases where the effects on well-being are reduced. If X is required to advise V2, Y has created the conditions

where X must actively serve the end of V2's education and career. In contrast, if X is not required to advise V2, and does not do so, V2 is not required to serve any particular end.

This consideration may have some force, but I doubt that it is normally sufficient to relieve X of the obligation in cases like *Advice*. The consideration alone does not seem very powerful. And it competes with the countervailing consideration that the burden that V2 will suffer if she is not advised is otherwise greater than the burden that X will suffer if he gives the advice.

I do not rule out the possibility that there are cases where wrongdoing makes a person's duties less stringent than those duties would be were the wrongdoer to have been unable to perform the relevant action. Indeed, I will explore two kinds of case below. But *The Infliction Argument* does not generally support a principle like *Fair Share II*, or a variation on that principle restricted to circumstances of low cost.

IV. Doing, Allowing, and Manipulation

In the rescue cases that we have just considered, our question was whether Y's wrongful action or inaction makes a difference to the stringency of X's positive obligations. Does it make a difference if Y's wrongdoing will affect whether X causes harm to others?

i) Doing and Allowing

Let us consider some variations on a familiar case of causing harm:

Trolley: A trolley is heading down a track towards five innocent people who will be killed if nothing is done. X can divert it to another track with one innocent person on it, who will then be killed.

Suppose that the *Doctrine of Doing and Allowing* permits turning the trolley, but would not permit X to turn it in:

Trolley 2: A trolley is heading down a track towards five innocent people who will be killed if nothing is done. X can divert it to another track with four people on it, who will then be killed.

X would act wrongly in *Trolley 2*, let us suppose, because although the saving of five lives is sufficiently important to justify killing one, the saving of five lives is insufficiently important to justify killing four.

Now compare:

Trolley 3: as *Trolley 2*, except Y wrongly tied the four to the track.

What difference does Y's wrongdoing make to the permissibility of X turning the trolley?

The objections to *The Responsibility Argument* above apply equally here. The fact that Y will be fully and primarily responsible for the deaths if X turns the trolley in *Trolley 3* need not have implications for X's responsibility if he does so, nor for the stringency of his duty not to turn it. To assess whether it does, we need to evaluate considerations independent of responsibility itself.

Now recall the argument that our duties can make us, or others, victims of wrongdoing and that this provides reasons for or against us having these duties. As the five and the four are strangers to X, it does not make much difference to X whether he has the duty or not. The main people who are affected by whether X has such a duty are the five and the four.

If X is not permitted to turn the trolley, the five will be victims of Y's wrongdoing. If X is permitted to turn it, the four will be victims of Y's wrongdoing. But if X is not permitted to turn the trolley, and does not do so, Y will have done something morally equivalent to preventing the five from being saved. In contrast, if X is permitted to turn it, and does so, Y will have done something equivalent to killing the four. It is plausibly worse to kill a person than it is to prevent the person from being saved. If so, the reason against X being permitted to turn the trolley is more powerful than the countervailing reason for his being permitted. Overall, then, if there is a difference of this kind between these cases, X's duty not to turn the trolley is even more stringent in *Trolley 3* than in *Trolley 2*.

ii) *Intentions and Manipulation*

We have been exploring the idea that the duties that a person has depends on how gravely different people, including the duty holder, would be wronged if that person has or lacks that duty. As there are many facts that can affect the gravity of wrongdoing, the implications of this idea will vary from case to case.¹⁰

For example, many people think that the gravity of a person's wrongdoing is affected by her intentions.¹¹ If so, the relevant intentions might also affect the duties that a person has. For example, suppose that in *Trolley 3* Y tied the four to the track in order to ensure that the five are not saved. If it is wrong for X to turn the trolley, Y will have intentionally prevented the five from being saved. If it is permissible for X to turn the trolley, and she does so, Y will not have killed the four intentionally. It might be a reason why X is permitted to turn the trolley that if X lacks this permission, and does not do so as a result, the five will be victims of Y's wrongful plan. I doubt that this consideration is sufficiently important to make a significant difference in cases of life or death, but it might have a role to play in cases of lesser harm.

¹⁰ We might also explore the possibility that impersonal considerations make a difference to what duties people have. For example, there might be a general and impersonal reason to thwart evil plans. For reasons of space, I leave this possibility aside.

¹¹ I am grateful to an associate editor of *Philosophy and Public Affairs* for encouraging me to explore this possibility.

A related idea might be important in the Gaza case introduced earlier. Those attacking Israel were alleged to have hidden in civilian areas in order to make it disproportionate for the Israelis to use force to defend themselves. Supporters of Israel's actions have argued that these intentions affect what it is proportionate to do.

Yishai Schwartz, for example, argues that Israel faced a paradox: 'Morality demands that Israel fight this war, but allows no way to fight it morally. In this conflict, reason itself seems to fail.'¹² Of course, this seeming paradox is easily dispelled – X can have a *pro tanto* duty to *a*, and yet if *a*ing would violate a stringent duty, it may be wrong all things considered for him to *a*. If I promise to take my kids to the zoo today, I have a *pro tanto* duty to do so, yet it is all things considered wrong to fulfill this duty if the only way to do it involves running someone over. Reason doesn't fail in this case. Similarly, Israel may have a *pro tanto* duty to fight the war, but if the only way to fight it is disproportionate, doing so is wrong.

Schwartz, though, finds a different solution to the 'paradox'. He writes:

There is, however, a way out of this paradox. And we find it at the moment we realize that Hamas' actions have made this war about more than Israel or Palestine; it's a war about future of morality in armed conflicts. For if Israel declines to fight, we live in a world where

¹² 'Israel's Deadly Invasion of Gaza is Morally Justified' *New Republic* July 21, 2014.

terror groups use their own civilians, and twist morality itself, to bind the hands of those who try to fight morally. In this world, cruelty is an advantage, and the moral are powerless in the face of aggression and indiscriminate attack. And make no mistake: The eyes of the world are on Hamas, and terrorist groups worldwide will—as they have for generations—learn from the tactics of Gazan terrorists and the world’s reaction. So if Israel allows Hamas’ human shields to defeat it now, we will all reap the results in the years to come.

But there is an alternative. We can say that there is a principle worth fighting and dying for: Civilians cannot be used to make just wars impossible and morality will not be used as a tool to disarm. And once we have that principle, the proportionality calculation changes. The deaths of innocents are not simply outweighed by Israelis’ right to live without daily rockets and terrorists tunneling into a kibbutz playground; but by the defense of a world in which terrorists cannot use morality to achieve victory over those who try to fight morally. It is the protection of that world, one in which moral soldiers still have a fighting chance, that justifies Israel’s operations against Hamas today. And it is that greater cause that decisively outweighs the terrible toll in innocent life.¹³

¹³ *Ibid.*

Schwartz' argument is unclear. Perhaps he intends to show that whilst the actions of the Israeli government seem disproportionate, they are part of an overall policy that is proportionate, and so are themselves proportionate. Suppose that Israel adopts a policy of discounting the moral significance of non-combatant deaths if militants use these non-combatants as human shields. Some might argue that executing this policy is permitted because doing so will discourage militants from hiding amongst civilian populations, and this will reduce the death toll overall. Thus the death toll involved in executing this policy is proportionate to the overall protection that it provides.

It is difficult to know whether this argument succeeds, but it is irrelevant to this paper: it relies on the incentive effects of responding to wrongdoing that I set aside in the introduction. I point to it only to distinguish it from a more fundamental argument that Schwartz seems attracted to: that the proportionality principle is itself sensitive to the intentions of wrongdoers.

He might intend something like:

Exploitation:

If:

- i) X has a *prima facie* or *pro tanto* duty to *a* to stop Y securing his wrongful aims; and
- ii) The harm that X will do to V as a side-effect if he *as* would make *ain*g disproportionate were no one else to have acted wrongly; and
- iii) Y puts V in harm's way with the intention of ensuring that X does not *a* with the intention of securing his wrongful aims; and
- iv) X complies with his duties.

Then:

- 1) If X is prohibited from *ain*g Y successfully prevents X from *ain*g.
- 2) Y then exploits the proportionality principle to ensure that X does not frustrate his wrongful aims.
- 3) There is a reason against the proportionality principle being vulnerable to being exploited to secure a person's wrongful aims.
- 4) Therefore the prohibition on X *ain*g is less stringent than it would be had V been in harm's way without Y having the relevant intention.

It does seem especially objectionable that people protect their ability to carry out wrongful attacks by intentionally exploiting the proportionality principle. Furthermore, it might be argued that there are especially strong reasons to frustrate their plan to evade preventive force by hiding amongst civilian populations. And it might be argued that the civilians themselves have a

reason to ensure that terrorists do not use them to evade preventive force. This might help to justify inflicting some extra cost on them.

Furthermore, if Y changes X's obligations by putting V in harm's way, Y controls X's conduct. X has an interest in not being controlled or manipulated in this way. Here is the basic idea. Y can influence X's conduct by altering her circumstances in a way that gives rise to her having duties. This often seems objectionable. It is even more objectionable, though, if the Y intends to control X's conduct.¹⁴ People are less vulnerable to manipulation if others cannot alter what it is permissible for them to do by making their actions harmful to others, if they do so with the intention of controlling them.

This argument has some force, but I doubt that it is sufficiently important to make any significant difference where life and death is at stake, such as the Gaza case itself. Suppose that the Israelis are permitted to strike. The non-combatants that they kill will then be the victims of the militants' wrongful and inept attempt to exploit the proportionality principle. They will have been harmed as a result of the militants' attempt to incorporate these non-combatants into their plans. These people will have suffered wrongful harm, or its equivalent, at the hands of those people wrongly hiding amongst them. They are not the intended targets of the wrongdoing, and that may make some difference to their claim not to be harmed, but this consideration

¹⁴ For further discussion, see V Tadros *Wrongs and Crimes* ch.5.

does not seem sufficiently powerful to make a significant difference to their right not to be killed.¹⁵

The argument may be important in cases where the harms inflicted are less severe, so there may be a modest exception to the more general idea that our duties not to harm others are more stringent when the harm that others will suffer is a consequence of wrongdoing.

V. Intervening Agency

While many will agree that X's duty not to kill is at least as stringent in *Trolley 3* as it is in *Trolley 2* many believe that intervening agency makes a difference to permissibility. They believe something like:

Intervening Agency: Other things equal, if X acting will result in V being harmed, the prohibition on X acting is less stringent if and because Y intervenes in a certain way between X acting and V being harmed.

As we will see, although this is not always clear from discussions of intervening agency, the most plausible version of this principle focuses on wrongful intervening agency in particular.

¹⁵ See, also J McMahan 'What Rights May be Defended by Means of War?' in C Fabre and S Lazar *The Morality of Defensive War* (Oxford: OUP, 2014) 151.

Many just war theorists think that harms caused by intervening agents, such as insurgents, count less in the proportionality calculation than those who are killed directly. Thomas Hurka canvasses a range of views, but is unsure which to accept.¹⁶ David Rodin and Jeff McMahan think intervening agency is significant in this way;¹⁷ Helen Frowe thinks that it is very significant.¹⁸ I argue that intervening agency is morally irrelevant.

Because war is complicated, let us consider cleaner pairs of cases. As we will see, it is difficult precisely to characterize intervening agency. But there is some intuitive contrast between cases of intervening agency and more standard cases where a person is placed in harm's way, such as those considered in the Section IV(i). Compare:

Trolley 3: A trolley is heading down a track towards five innocent people who will be killed if nothing is done. X can divert it to another track with four people on it who will then be killed. Y wrongly tied these four to the track.

Trolley 4: A trolley is heading down a track towards five innocent people who will be killed if nothing is done. X can divert it towards

¹⁶ See 'Proportionality in the Morality of War' (2005) 33 *Philosophy and Public Affairs* 34, 47-50.

¹⁷ D Rodin 'The Myth of National Self-Defence' in Fabre and Lazar *The Morality of Defensive War* 82; McMahan 'Proportionate Defense', 36.

¹⁸ *Defensive Killing* (Oxford: OUP, 2014) ch.5.

another track with four people on it. The trolley will stop on the track before it reaches the four but Y will immediately restart the trolley, killing them. If X does not act, the four will be unharmed.

In both *Trolley 3* and *Trolley 4*, if X is permitted to turn the trolley, the four will be killed, and they will be victims of wrongdoing. But it seems that they are the victims of wrongful intervening agency only in *Trolley 4*.

Friends of *Intervening Agency* are not very clear about what they mean by intervening agency. Views about intervening agency have two dimensions, concerning the relevant kind of agent and the nature of the intervention.

i) *What Kind of Agency?*

What deontic and agential properties make Y an intervening agent on the best version of *Intervening Agency*? Intervening agents can either act permissibly or wrongly, and they can be non-responsible, non-culpably mistaken, or both responsible and culpable.

First let us focus on intervening agents who fulfill standard conditions of responsibility, but who act permissibly. Consider:

Trolley 5: As *Trolley 4*, except Y, a fully responsible agent, can save one hundred people by pressing a button that will also restart the trolley. Pressing this button will kill the four. Had X not acted, Y would have

been able to save the one hundred without killing anyone. If X diverts the trolley, Y will kill the four.

Y's intervention seems permissible. The fact that he acts to kill the four cannot plausibly make X's act easier to justify than it is in *Trolley 3*. Creating the conditions where one person justifiably harms another seems no different to causing harm directly.

Here is why. By turning the trolley, X makes it inevitable either that the four die, or that Y does not save the one hundred. The fact that it is permissible for Y to save the one hundred by killing the four suggests that it is at least as important that the one hundred are saved as that the four are not killed. X's act thus gives Y no choice but to kill the four, or to do something at least as morally significant - in this case, not saving the one hundred. The fact that Y's agency determines which of these bad results occurs cannot make it permissible for X to act. This argument does not apply where Y acts wrongly.

Now focus on the agential qualities of the intervener. An agent is a being that can act. Not all such beings are responsible - they are not responsible if they lack certain capacities. It is disputed what these capacities are, but non-human animals and very small children are clear examples of non-responsible agents - they act but they are not responsible for these acts. It is hard to believe that non-responsible intervening agents make a difference to permissibility - make Y a non-human animal, or a very small child, and there is no plausible difference between *Trolley 3* and *Trolley 4*. The same thing

is true of intervening agents who make reasonable mistakes – these people are responsible for their acts, but they do not anticipate the consequences of them for good reason.¹⁹ In *Trolley 4*, if Y does not know that the four are further on down the track, and X knows Y will restart the trolley as a result, Y's intervening act cannot plausibly affect the permissibility of X's act.

The most plausible version of *Intervening Agency* claims that Y's intervening agency makes a difference to the permissibility of X's act only if Y will be culpable for acting wrongly. This is the view I will consider.

ii) *What is an Intervention?*

It is much trickier to find the best account of intervention. We know that it is insufficient for Y to be an intervening agent that the harm that V suffers comes about because of Y's act. This was true of the cases we considered in the last section, such as *Trolley 3*, but these were not cases of intervening agency. Another possibility is that Y is an intervening agent only if X's act results in harm because of Y's act and Y's act occurs after X's act. I don't think

¹⁹ A more plausible view, that a reviewer for *Philosophy and Public Affairs* suggested, is that intervening agency makes a difference in cases where the intervening agent is aware of the harm that she will cause, but makes a reasonable mistake about the liability of those who will be harmed. If intervening agency makes no difference between what a person owes to the victims in cases of wrongful harm, as I argue below, it is hard to believe that it makes a difference to what the victims are owed in this case, so I set this view aside.

that whether a person is an intervening agency is typically understood simply as a temporal matter though.

Consider:

Drop: As Trolley 3, except Y drops the four onto the track after X pulls the lever.

I don't think that friends of *Intervening Agency* understand Y as an intervening agent in this case, though it is difficult to be sure.

I think that those who think intervening agency important intend to distinguish cases where Y causes V to be in harm's way, either before or after X acts, and cases where the chain of events that leads to V being harmed 'goes through' Y's act. We might, for example, compare a person who puts a domino in a line of dominos the last of which causes harm when it falls and the person who herself stands in the line of dominos and pushes the next one when the last one touches her. Only the latter's act is 'in the chain' in the relevant sense. It is very difficult to make the idea of an act being 'in the chain' more precise. As I will reject the significance of intervening agency even in central cases, I need not characterize what an intervening agent is more precisely.

iii) *Doubts about Intervening Agency*

Friends of *Intervening Agency* typically appeal to the idea that Y's responsibility for the harm that V suffers diminishes X's responsibility for that harm if Y is an intervening agent. As we saw when considering *Demolition 2*, two people can be fully responsible for causing harm even when they are not working together. In those cases, Y's responsibility in no way diminishes X's responsibility. Why should Y's responsibility have any more effect on X's responsibility where Y is an intervening agent? Furthermore, considerations of responsibility don't seem to support *Intervening Agency* in particular. Such considerations seem equally relevant in *Demolition 2*, and that is not a case of intervening agency.

Some might argue that respect for intervening agents should lead us to discount the effects of their conduct. We should presume that others will not act wrongly, even when we predict that they will. This view would have implications not only in cases of intervening agency, as I think this idea is commonly understood, but also in cases such as *Drop*, where the agent is yet to create the circumstances where one's act will result in harm.

As I have addressed this issue elsewhere, I will keep my response brief.²⁰ The view that we should act on the presumption that others will not act wrongly is often defended on the grounds that were we not to do this we would treat the person as incapable of acting permissibly. Some also claim

²⁰ V Tadros 'Controlling Risk' in A Ashworth, P Tomlin and L Zedner *Prevention and the Limits of the Criminal Law* (Oxford: OUP, 2013).

that acting on the prediction that a person will act wrongly treats her like a child.

These claims are false. Making a prediction about the behavior of others implies nothing about what she is capable of. Compare the fact that that predicting that a person will act rightly does not imply that she could not act wrongly, and the fact that predicting a person will act rationally does not imply that she could not act irrationally. It is equally untrue that predicting that a person will act wrongly implies that the person cannot refrain from acting wrongly.

Furthermore, people often make predictions about what they themselves will do amongst the range of actions they are capable of doing. When a person predicts that she will act wrongly, she does not treat herself as incapable of acting permissibly, nor does she treat herself like a child. Others no more treat her in these ways when they act on similar predictions.

Finally, the idea that we should do as though people will not act wrongly even when we predict that they will has counterintuitive implications for the use of state resources – we would have stronger reasons to divert those resources to the prevention of natural disaster than the prevention of future harmful wrongdoing through preventative policing, which is not very plausible. This is especially implausible when a person is about to act wrongly: would it really be disrespectful to intervene to prevent

an assassin from shooting his victim on the grounds that he might do the right thing and desist?²¹

Perhaps it might be argued that it is morally basic that we are more responsible for the harms that we directly cause than the harms that we indirectly cause. In *Trolley 3*, X directly harms the four by turning the trolley. In *Trolley 4*, X indirectly harms the four through Y's action. There is something intuitive about this idea, but it is difficult to clarify it, let alone to defend it. The intuitive difference is that in *Trolley 4*, the consequences of X's action do not occur 'through' Y's action. It might be argued that a causal chain is weakened or broken by wrongful acts. There is such a weakening or break in the causal chain in *Trolley 4*, but not in *Trolley 3*.

But we have already seen intervening acts are only plausibly relevant to permissibility if they are culpable and wrongful. It is difficult to see how either culpability or wrongdoing affect whether one event causes another. And the idea that intervening agency is morally basic is not in itself intuitive. Where two acts, X's and Y's, are part of a causal chain leading to harm to others, and X's act will also save others from being harmed, why should whether X's act or Y's act occurs earlier in the chain affect permissibility?

To reinforce this concern, consider:

²¹ I am grateful, here, to an anonymous reviewer for *Philosophy and Public Affairs*.

Prime and Launch: X and Y are in separate booths. Each has an identical button in front of him. If X and Y both press their buttons in any order, a weapon will be primed and launched, which will kill four people. The first press, whichever button it is, primes the weapon. The second, whichever it is, launches the weapon. X can rescue five people, but as a side-effect he will press his button. He knows that Y will certainly press his button either immediately before X does so, or immediately afterwards. He does not know which.

If X presses before Y, Y is an intervening agent. Y will directly launch the missile killing the four. If X presses after Y, Y is not an intervening agent. X will directly launch the missile killing the four. *Intervening Agency* supports the view that X has a reason to find out whether Y has pressed to determine whether he is permitted to save the five, for Y's failure to press militates in favour of X saving the five. But X does not seem to have a reason to find out whether Y has pressed.

Furthermore, suppose that X can save more lives if he presses later than if he presses earlier. He is certain, either way, that the weapon will be launched at the four. *Intervening Agency* implies that X has a reason to press now rather than later, even though he will save fewer lives if he does so because if he waits, Y is less likely to be an intervening agent, and hence it is more likely that X will act wrongly. This is also hard to believe.²² *Prime and*

²² I am grateful, here, to Jeff McMahan.

Launch thus supplies powerful support for the irrelevance of intervening agency to permissibility.

Perhaps it might be argued that *Prime and Launch* does not have implications for intervening agency more generally. It is somewhat intuitive that X and Y together launch the weapon. This may be because they both contribute to a system that is designed to launch the weapon. But it is hard to see why it should matter to permissibility whether there is such a system or not.

iv) *The Harmful Effects of Self-Defence*

Perhaps we should reject the intuitive view that there is no difference between pressing first and pressing second in *Prime and Launch* in the light of a wider range of cases. Consider:

Defence and Harm: X attacks Y. Y knows that if he successfully defends himself against X, Z, who is a friend of X's will be enraged and kill V.

Some believe that X is permitted to fend off Y's attack, and this view seems intuitive, even though one would not normally be permitted to save one's life

in a way that will result in an innocent person being killed as a side-effect. It has been argued that intervening agency explains this intuition.²³

This case is not ideal for investigating the significance of intervening agency in particular for a number of reasons. First, it is not clear what kind of case we ought to compare *Defence and Harm* with to consider whether intervening agency, in particular, explains why Y seems permitted to defend himself. Without a clear comparison, which isolates intervening agency in particular, we should not be confident about the source of our intuitions.

Second, in *Defence and Harm* we are concerned with the permissibility of a person protecting her own interests through self-defence in a way that affects others. This is in contrast with cases like *Trolley 4* and *Prime and Launch*, where we are not concerned with the interests of the person whose duties we are considering. We could eliminate this feature of the case as follows. Suppose that a stranger, S, can prevent Y from being killed. If she does so, Z will kill V. S does not seem permitted to do this. This suggests that *Intervening Agency* does not explain our intuitions in *Defence and Harm*, as Z is still an intervening agent in this variation.

Third, remaining passive in the course of an attack when one is able to resist is extremely difficult, and may make the attack much more harmful and humiliating. If Y fends X off, another murder will occur, but no one is required to refrain from struggling against his attacker when he can do so.

²³ For discussion of similar cases, drawing this conclusion, see H Frowe *Defensive Killing* (Oxford: OUP, 2014) 129-31; McMahan 'What Rights May be Defended by Means of War?' 150-1.

And we may have interests in a permission to fight back that are independent of our interests in succeeding. For example, by fighting back we express our objection to what is happening to us.²⁴ This helps to explain why we are inclined to think that Y is permitted to defend himself against X.

These ideas might not all be morally important, but they are sufficient to cast doubt on the idea that *Defence and Harm* supports *Intervening Agency*.

v) *Gandhi: Not Non-Violent After All*

A further troubling case was noted in the Introduction. There I asked whether Gandhi's protests were really non-violent in the morally relevant sense, because of the harm they predictably caused. Many people will find it powerfully intuitive that Gandhi acted permissibly, and that the reason why has something to do with the fact that the colonial authorities are responsible for the harms inflicted on the protesters, and this affects Gandhi's responsibility for the harm that results from his actions. The view that this is because they are intervening agents is initially tempting.

When we consider pairs of cases cleanly, though, intervening agency does not seem relevant in this case either. Consider:

²⁴ See, especially, D Statman 'On the Success Condition for Legitimate Self-Defence' (2008) 118 *Ethics* 659.

Gandhi I: Gandhi organizes a non-violent protest to secure Indian independence along some cliffs. Doing so causes a landslide, killing V.

Gandhi II: Gandhi organizes a non-violent protest in town to secure Indian independence. Y wrongly takes a dangerous route along the cliffs to avoid the protest. This causes a landslide, killing V.

I doubt that these cases are importantly different. Yet in the latter case, but not the former, an intervening agent causes V's death.

Perhaps the role of the police makes us more inclined to distance Gandhi from the harm that they caused as a result of his protests. This might be thought especially true because the police intend to stop Gandhi protesting peacefully. As I suggested when discussing the Gaza case, we find it difficult to accept that wrongdoers can prevent us from preventing their wrongdoing by making our pursuit of this aim disproportionate. But this idea has nothing to do with intervening agency as such. Wrongdoers might or might not be intervening agents, and yet have equal intentions to exploit morality to restrict actions that would otherwise be commendable.

Compare:

Gandhi III: Gandhi organizes a non-violent protest to secure Indian Independence along some cliffs. A police officer credibly threatens him

that if the protest goes ahead, he will start a landslide, killing V, who has been tied to a tree at the bottom of the cliffs.

Gandhi IV: As *Gandhi III* except the police officer ties V to a tree at the bottom of the cliffs so that Gandhi's protest, if it goes ahead, will directly kill V.

Suppose that the police officer in both cases aims to prevent the protests from going ahead by putting V in harm's way. He will be an intervening agent in *Gandhi III* but not *Gandhi IV*. This does not seem to make a difference to the permissibility of the protest going ahead.

This judgement can be reinforced as follows. If intervening agency reduces the significance of harm, we have a reason to act in a way that will harm others through intervening agency rather than harming others directly, even when wrongdoers will be responsible for the harm in both cases. This, though, does not seem very plausible. Consider:

Gandhi V: Gandhi organizes a protest to secure Indian independence along one of two cliffs. If he takes the east cliff, a police officer will start a landslide, killing V, who has been tied to a tree at the bottom of the cliff. If he takes the west cliff, he will directly start a landslide. This will kill W, who has been tied to a tree at the bottom of the cliff by the same police officer.

Intervening Agency implausibly implies that Gandhi has a reason to choose the east cliff rather than the west cliff.

Although it seems, at first sight, attractive to believe that Gandhi is not responsible for the harm done to victims of police brutality, this view does not survive scrutiny. This is not to say that Gandhi acted wrongly all things considered, of course. That depends on whether the harms that resulted from his actions were disproportionate to the good that he secured through them. It also depends on the fact that most of those who were harmed through police brutality were participants in his peaceful protests, and these people voluntarily assumed the risk of being harmed. These facts, rather than the fact that the police officers were intervening agents, better support the view that his protests were justified.

Conclusions

Many people think that harm counts less against our actions if that harm results from the wrongdoing of others. Though this view has a great deal of initial intuitive appeal it is not generally true. In Section IV(ii) We saw some support for some narrower principles permitting us to fail to prevent harm, or to harm others. For example, we may be permitted to cause, or fail to prevent, harm to others because otherwise we will be the intended victims of wrongdoing. And we may be permitted to do so to avoid others exploiting

the proportionality principle to secure their unjust aims. But these ideas do not support principles such as *Fair Shares II* that grant a more general permission not to take up the slack left by others. We found no support for the view that the responsibility of others for the harm that results from our conduct affects our duties, whether or not the person responsible acts or omits, and whether or not the person responsible is an intervening agent. And we found some support for the unpopular view that our duties to prevent harm, and not to cause harm, are more stringent when that harm will occur as a result of the wrongdoing of others.