

# Should We Prevent Optimific Wrongs?

ANDREAS L. MOGENSEN

*Jesus College, Oxford*

Most people believe that some optimific acts are wrong. Since we are not permitted to perform wrong acts, we are not permitted to carry out optimific wrongs. Does the moral relevance of the distinction between action and omission nonetheless permit us to allow others to carry them out? I show that there exists a plausible argument supporting the conclusion that it does. To resist my argument, we would have to endorse a principle according to which, for any wrong action, there is some reason to prevent that action over and above those reasons associated with preventing harm to its victim(s). I argue that it would be a mistake to value the prevention of wrong acts in the way required to resist my argument.

## I

According to act-consequentialism, the ends always justify the means: any action is morally permissible if it brings about an outcome that is the best available, considered from an agent-neutral perspective.<sup>1</sup> In the jargon of moral philosophy, actions of this kind are called *optimific*. Most people reject the view that optimific actions are always permissible. In some cases, they think, bringing about the outcome that is best from an agent-neutral perspective is morally wrong. Here is one paradigmatic example:<sup>2</sup>

### *Footbridge*

A runaway trolley is on course to hit five people whom it will crush to death. You are on a footbridge overhead. In order to prevent the five from dying, you can push a hiker wearing a heavy backpack into the

<sup>1</sup> Some philosophers put forward views that are like act-consequentialism except that they permit an agent-relative ranking of outcomes. See Jamie Dreier, 'Structures of Normative Theories', *The Monist* 76 (1993), pp. 22–40; Dreier, 'In Defence of Consequentializing', *Oxford Studies in Normative Ethics* 1 (2011), pp. 97–119; Douglas Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford, 2011). For the purposes of this article, 'act-consequentialism' should be understood to require an agent-neutral ranking, such that the ordering of outcomes as better or worse does not vary from agent to agent.

<sup>2</sup> For the sake of realism, this case is modified from its canonical description in Judith J. Thomson, 'The Trolley Problem', *Yale Law Review* 95 (1985), pp. 1395–415. I encountered this improved version in a talk by Eric Schwitzgebel.

path of the trolley. Their combined weight will bring the machine to a halt, but the hiker will be crushed to death. There is no other way to save the five.

Most people believe that it would be wrong to sacrifice the hiker for the sake of saving the five in *Footbridge*.<sup>3</sup> Non-consequentialist theories are built around intuitions like this.

In this article, I'll assume that these intuitions are correct. We have settled, then, how I should act if I find myself in *Footbridge* and in other cases like it. The question I want to address concerns how I should react to actions that third parties might undertake in such contexts. More exactly, I want to consider whether we should prevent others from carrying out optimistic wrongs. Consider the following scenario:

*Footbridge\**

A runaway trolley is on course to hit five people whom it will crush to death. We are on a footbridge overhead. In order to prevent the five from dying, you or I might push a hiker wearing a heavy backpack into the path of the trolley. Their combined weight will bring the machine to a halt, but the hiker will be crushed to death. There is no other way to save the five. I see that you are just about to push the hiker onto the tracks.

Should I prevent you from doing so?

The question of what I ought to do in *Footbridge\** would appear to be open even if we've settled what I should (not) do in *Footbridge*. After all, one of the most widely endorsed non-consequentialist moral principles is:

*The Doctrine of Doing and Allowing (DDA)*<sup>4</sup>

We have greater reason not to actively bring about harm than to allow harm by omission.

The following cases illustrate the intuitive force of the distinction<sup>5</sup>

<sup>3</sup> A large web-based survey found that only 11 per cent of participants thought it permissible to push the person onto the tracks in this kind of case. See Marc Hauser, Fiery Cushman, Liane Young, R. Kang-Xing Jin, and John Mikhail, 'A Dissociation Between Moral Judgments and Justifications', *Mind & Language* 22 (2007), pp. 1–21.

<sup>4</sup> For a summary of recent research on DDA, see Fiona Woollard, 'The Doctrine of Doing and Allowing', *Philosophy Compass* 7 (2012), pp. 448–69.

<sup>5</sup> These cases are due to Frances Kamm, *Morality, Mortality*, vol. 2 (Oxford, 1996), p. 90. They are based on examples first described by Philippa Foot in 'Killing and Letting Die', reprinted in her *Moral Dilemmas* (Oxford, 2002), pp. 78–87.

*Ambulance*

You are driving an ambulance with five injured people. In order to get them to the hospital on time and prevent their deaths, you must drive over a pedestrian who blocks your path, killing her.

*Ambulance\**

You are driving an ambulance with five injured people. In order to get them to the hospital on time and prevent their deaths, you must drive past an injured pedestrian who will die unless you stop to help her.

Intuitively, it is permissible to allow the pedestrian to die in *Ambulance\** although you may not actively bring about her death in *Ambulance*. We may wonder whether DDA permits us to treat *Footbridge* and *Footbridge\** similarly: that is, whether the distinction between action and omission permits us to allow the hiker to be pushed, although we may not push her. More generally, we may wonder whether the outcomes associated with optimific wrongs are such that we can allow them, though we may not actively bring them about.

Some may feel that DDA has little force here. They might think that *Footbridge\** is simply not a case where the distinction between doing and allowing makes such a difference.<sup>6</sup> Intuitions to the effect that one ought to intervene to prevent optimific wrongs are reported by Frances Kamm and Jeff McMahan.<sup>7</sup> Whilst not immune to the pull of the intuition that one ought to intervene in *Footbridge\** and in other cases of its kind, I want to challenge that conclusion in this article. I will show that there exists a plausible argument for the contrary position.

If we came to believe that it is permissible to allow optimific wrongs, this would have important consequences for ethical theory. On the one hand, it would impact our understanding of why actions of this kind are wrong. A question exists whether non-consequentialist constraints are *agent-centred* or *victim-centred*. That is: if I may not push the hiker in *Footbridge*, is this because I have reason to avoid *my being involved in this way* in her dying or is it because of some fact to do with *her dying in this way*, quite apart from any consideration of my involvement? As Kamm notes, questions of how we should address attempts by others to carry out optimific wrongs provide a natural test-case for this issue: the view that we should prevent optimific wrongs

<sup>6</sup> It should be uncontroversial for those who accept DDA that it is worse for me to push the hiker than to allow her to be pushed. The question is whether the difference is so great that permitting her to be pushed is permissible, and not merely wrong to a lesser extent.

<sup>7</sup> See Frances Kamm, 'Rights beyond Interests', in her *Intricate Ethics* (Oxford, 2007), pp. 237–84, at 252, and Jeff McMahan, 'Intention, Permissibility, Terrorism and War', *Philosophical Perspectives* 23 (2009), pp. 345–72, at 350.

naturally supports the view that non-consequentialist constraints are victim-centred, whereas the view that we can allow optimific wrongs accords with the view that such constraints are agent-centred.<sup>8</sup>

More importantly, the view that we should allow optimific wrongs appears to have significant implications when it comes to issues in applied ethics. It might be taken to imply, for example, that even if embryo experimentation is morally wrong because it involves killing human beings,<sup>9</sup> you and I should not try to end such experiments, provided that continued experimentation is recommended on consequentialist grounds. Similarly, if experimenting on animals for medical purposes is morally wrong because animals have rights,<sup>10</sup> this would not suffice to show that we should seek to end such experiments, unless a consequentialist defence of animal experimentation is also unavailable.

Here is the structure of this article. In the next section, I will set out the argument challenging the view that we ought to prevent optimific wrongs. I identify two moral principles that could be proposed to resist this argument: the *Wrong-Preventing Principle* and the *Scope Principle*. I dismiss the *Scope Principle* on the grounds that it sanctifies a cognitive bias; I concentrate my attention on the *Wrong-Preventing Principle*. In section III, I provide three arguments against the suggestion that we should resist my argument from section II by appeal to the *Wrong-Preventing Principle*, pointing to various unpalatable implications that arise from doing so. I conclude that my argument is sound.

## II

As stated, my view is that there exists a plausible argument against the view that we should prevent optimific wrongs. The argument is straightforward, although it makes reference to a hypothetical example (called *Tracks*) that I have yet to introduce. I will give the argument first and then present the case. The argument runs as follows:

1. If we should prevent optimific wrongs, I should intervene in *Footbridge\**.
2. If I should intervene in *Footbridge\**, I should intervene in *Tracks*.

<sup>8</sup> Kamm, 'Rights beyond Interests'.

<sup>9</sup> For this view on the moral status of embryos, see Robert P. George and Alfonso Gomez-Lobo, 'Statement of Professor George (Joined by Dr. Gomez-Lobo)', in his *Human Cloning and Human Dignity: An Ethical Inquiry* (Washington, D.C.: The President's Council on Bioethics, 2002), pp. 258–65; Gomez-Lobo, 'The Moral Status of the Human Embryo', *Perspective in Biology and Medicine* 48 (2005), pp. 201–10.

<sup>10</sup> For this view on the moral status of animals, see Gary Francione, *Animals as Persons: Essays on the Abolition of Animal Exploitation* (New York, 2008); Tom Regan, *The Case for Animal Rights* (Berkeley, 1983).

3. It is not the case that I should intervene in *Tracks*.

*Therefore*, it is not the case that we should prevent optimific wrongs.

Even without knowing about the case I call *Tracks*, we can already note two important facts about this argument: that premise 1 is clearly true, and that the argument is straightforwardly valid. Therefore, the soundness of the argument rests on the truth of premises 2 and 3. To assess their truth, we need to get acquainted with the case I call *Tracks*. Here it is:

*Tracks*

You discover a hiker with a large bag lying on the tracks. A runaway trolley is coming towards her and she will die in the event of a collision, thereby stopping the trolley. Your immediate inclination is to pull her to safety. However, you discover that doing so will allow the trolley to proceed on its path and run over five people who are located further up the track.

In this case, it seems false to say that what you ought to do is intervene and pull the hiker to safety.<sup>11</sup> In fact, it seems to me that it would be seriously wrong to intervene in this case because of the manner in which this would implicate you in the death of the five. Here, the needs of the many outweigh the needs of the few. Thus, premise 3 in my argument seems true. The soundness of the argument hinges, then, on premise 2.

This premise rests on the idea that *Footbridge\** and *Tracks* are sufficiently similar that if I ought to behave one way in the one case, I should behave that way in the other too. It should be obvious that *Footbridge\** and *Tracks* are very similar. Both involve trolleys, tracks and hikers; both involve the same number of potential victims; both involve the possibility of allowing someone to be used as a means to save a greater number. It is perhaps not so obvious that these two cases are similar in *all* relevant respects. So far as I can tell, there are two – and only two – differences between the cases to which one could appeal to support the view that one must act to save the hiker in the *Footbridge\** but not in *Tracks*.

First, one might appeal to the fact that in *Footbridge\**, not only would I save the hiker, but I would also prevent the action that would result in his death. Since that action is wrong, we might attach some intrinsic importance to ensuring that it does not occur: we may think it valuable

<sup>11</sup> Shelly Kagan offers a similar case, about which he draws the same conclusions in *The Limits of Morality* (Oxford, 1989), p. 164.

in and of itself that a wrong action is thwarted. Larry Temkin claims that ‘acting rightly is *itself* a good-making feature, and acting wrongly *itself* a bad-making feature, of the outcomes of which they are a part.’<sup>12</sup> Thus, there may be a reason to intervene in *Footbridge\** that does not apply in *Tracks*: a reason associated not with ensuring the survival of the hiker, but with preventing a murder. We might capture this proposal via the following principle:

*The Wrong-Preventing Principle*

For any wrong action, there is some reason to prevent that action over and above those reasons associated with preventing harm to the victim(s).

There is a second potential difference between *Footbridge\** and *Tracks* to which one might appeal, which has to do with the question of who is and is not already threatened. In *Footbridge\**, when one arrives on the scene, the hiker is not yet in the line of the trolley, whereas the five are. In some sense, the hiker is presently safe and they are not. In *Tracks*, the five are not threatened by the trolley owing to the presence of the hiker; the hiker is found in the line of the trolley. Here, the five are safe and the hiker is not. Someone might believe that these facts concerning the prior distribution of the threat allow one to remain passive in *Tracks* but not in *Footbridge\**. This proposal is reminiscent of a view set out by Judith Thomson and James Montmarquet in discussing permissible harms.<sup>13</sup> On their view, it is crucial to the question of whether we can permissibly harm some in benefiting a greater number whether the person who might be harmed is, in some sense, within the scope of an existing threat: if they are within its scope, harming them may be justified in cases where it would otherwise be impermissible. The hiker is outside of the scope of the trolley in *Footbridge\**, but is within its scope in *Tracks*. We might therefore try to justify saving the hiker in the *Footbridge\** but not in *Tracks* by adopting the following principle:

*The Scope Principle*

We have greater reason to save someone from being killed by some potential threat if she is not yet within the scope of the threat that would kill her.

These principles appear to me to exhaust our options in so far as we mean to reject premise 2 in my argument: there is no other

<sup>12</sup> Larry Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford, 2012), p. 205.

<sup>13</sup> See Judith J. Thomson, ‘The Trolley Problem’, and James Montmarquet, ‘On Doing Good: the Right and the Wrong Way’, *Journal of Philosophy* 79 (1982), pp. 439–55.

plausible way in which to explain why one might be required to intervene in *Footbridge*\* but not in *Tracks*. However, I will not give these two principles equal treatment. I believe the *Scope Principle* can be dismissed rather quickly. It would be misguided to endorse this principle, I believe, because if our intuitions are responsive to facts about who is and is not already threatened, those intuitions reflect *loss aversion*.<sup>14</sup>

Loss aversion must be understood against the background of *prospect theory*, a model of decision-making developed by Kahneman and Tversky.<sup>15</sup> According to prospect theory, decision-making involves the adoption of some state of affairs to represent the neutral point, with various possible outcomes represented as gains or losses therefrom. 'Loss aversion' refers to our tendency to treat losses as especially bad: failure to achieve a gain of some magnitude is not regarded as being as bad as a loss of the same magnitude.

It is well known that loss aversion can bias people's judgements about saving lives.<sup>16</sup> Kahneman and Tversky asked subjects to imagine that an unusual Asian disease might kill 600 people but could be mitigated by one of two programmes. If programme A is chosen, it is certain that 400 will die and 200 will be saved. If programme B is chosen, there is a one-third probability that all 600 are saved and a two-thirds probability that all die. Most people prefer programme A to programme B if the outcome of programme A is described as being simply '200 saved'. However, most prefer B to A if the outcome of A is described as '400 dead'. Given 600 potential victims, these descriptions are equivalent, but they lead people to think about the potential outcomes in different ways. Consider the potential outcome in which programme B is chosen but no one survives. By focusing on those who would survive under programme A, the first description ('200 saved') leads people to treat this potential outcome from programme B as involving a loss of 200 lives, whereas the second description ('400 dead') leads people to treat this outcome as simply a failure to realize a potential gain of 200 lives saved. Owing to loss aversion, losses are treated as worse than failures

<sup>14</sup> Tamara Horowitz argues that intuitions taken to support DDA reflect nothing more than loss aversion in 'Philosophical Intuitions and Psychological Research', *Ethics* 108 (1998), pp. 367-85. However, Kamm argues convincingly that we continue to regard killing as worse than letting die even in cases where letting die represents a loss and/or killing represents a gain foregone: see 'Moral Intuitions, Cognitive Psychology, and the Harming/Not-Aiding Distinction', in *Intricate Ethics*, pp. 422-9.

<sup>15</sup> Daniel Kahneman and Amos Tversky, 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica* 47 (1979), pp. 263-92.

<sup>16</sup> Daniel Kahneman and Amos Tversky, 'The Framing of Decisions and the Psychology of Choice', *Science* 211 (1981), pp. 453-8.

to achieve corresponding gains, and so the choice of description alters people's risk-tolerance and their preference amongst the programmes.

As this case makes clear, moral intuitions that reflect loss aversion are readily susceptible to biasing effects that affect the choice of status quo, leading us to adopt inconsistent prescriptions in the face of one and the same choice situation. For this reason, it appears mistaken to lend epistemic weight to intuitions known to reflect loss aversion. In so far as our intuitions are responsive to facts about who is and is not already threatened, it appears highly plausible that those intuitions *do* reflect loss aversion. Because loss aversion is a bias, we should resist any inclination to attach moral importance to this difference.

We should dismiss the *Scope Principle*, therefore, leaving the *Wrong-Preventing Principle* as the only plausible explanation for why one ought to intervene in *Footbridge\** but not in *Tracks*. I believe that the *Wrong-Preventing Principle* has greater *prima facie* plausibility in any case. To secure the case for premise 2 in my argument, I need then to argue against appealing to this principle in order to support the verdict that one should react differently in *Footbridge\** than in *Tracks*.

### III

In this section, I will present three arguments for why we should not appeal to the *Wrong-Preventing Principle* to resist premise 2, but rather adopt the conclusion of my argument instead.

Here is the first argument. The world that we inhabit is far from utopian. In many places, especially in developing countries, people are very badly off through no fault of their own. Ideally, we would like to help them all. Unfortunately, our time and money is limited, so we are forced to be selective. Of the evils that abound, some are the result of unjust actions perpetrated by others, whereas others are naturally occurring. Note, then, that if we accept the *Wrong-Preventing Principle*, we should be required to give priority to causes that eliminate injustices rather than natural evils.

Of itself, this may not sound too bad. People do appear to attach particular importance to eliminating wrongs. The names 'Martin Luther King Jr' and 'Nelson Mandela' are justly famous, but far fewer recognize 'Jonas Salk' or 'Norman Borlaug'. This suggests that we give some kind of priority of importance to the righting of wrongs. Whether we are right to do so is, of course, a different matter. It is easy to get the sense that we place too much emphasis on the elimination of injustice and do not focus enough on overcoming the evils of nature. In any case, I believe that the *Wrong-Preventing Principle* can be shown to have intuitively unacceptable implications because of the manner in which it discounts the value of human welfare.



Here is why. If we endorse the *Wrong-Preventing Principle*, this means that in some cases we should choose a course of action that provides a lesser benefit because it also ensures the prevention of a wrong: we leave people worse off for the sake of there being fewer wrong acts. This applies even with a single set of potential beneficiaries: if a group of people is afflicted by harms arising from some natural phenomenon and another kind of harm arising from wrongful actions, the *Wrong-Preventing Principle* may require that we prevent the wrongful acts rather than mitigating the natural phenomenon even if this means making the beneficiaries all-things-considered worse off. The *Wrong-Preventing Principle* may imply, for example, that it is more important to improve the economic rights of women in developing countries than to protect the same women against malaria, even if the women would be significantly better off (now and in future) as a result of our attempts to reduce the incidence of malaria. Everything hinges, of course, on what weight we attach to eliminating wrong acts. However, we already have some reason to believe that the weight should be considerable. Suppose that the *Wrong-Preventing Principle* does in fact explain why one should intervene in *Footbridge\** although one should not intervene in *Tracks*. If one follows these prescriptions, the fatality rate in *Footbridge\** comes out as five times higher than in *Tracks*. This suggests that the moral importance of preventing wrong acts is sufficient to outweigh at least a fivefold increase in harm. However, this seems quite implausible. It is surely a misguided fetishism to sacrifice human welfare to such an extent for the sake of there being fewer wrong acts.

My second and third arguments against the *Wrong-Preventing Principle* are similar in certain respects. Both will highlight the manner in which anyone who believes that the *Wrong-Preventing Principle* is the key to explaining why we should prevent optimific wrongs turns out to be committed to allowing people to be harmed and killed in various ways that do not seem, intuitively, to be any less worthy of prevention than the potential action described in *Footbridge\**. In asking us to intervene in *Footbridge\** but not in these other cases that I'll discuss, the *Wrong-Preventing Principle* commits us to drawing moral distinctions where there appear to be none.

Here, then, is the second argument. Understood as sufficient of itself to explain the difference between *Footbridge\** and *Tracks*, the *Wrong-Preventing Principle* implies that one has no obligation to intervene if *Footbridge\** is redescribed in such a way that the pushing of the hiker would be carried out by something other than a person. For example, the hiker might be pushed by a dog or by some kind of mechanical apparatus. Since this involves no wrongful action, I no longer have decisive reason to intervene and prevent the pushing.

This implication may be enough to make some uncomfortable. We can add to our discomfort by highlighting that such possibilities are not quite as fanciful as they may seem. We live in an age of increasing automation. Numerous self-driving vehicles already exist and others are due to appear within the near future. As has been pointed out, it may become necessary for self-driving vehicles to take account of morally relevant information and ‘decide’ certain ‘moral dilemmas’.<sup>17</sup> For example, an autonomous vehicle carrying someone who is injured might drive at speeds that would ordinarily be too dangerous. Although autonomous vehicles may respond to morally relevant considerations, I assume that those AIs likely to appear within the next decade are sufficiently primitive that they would be inappropriate as targets for moral responsibility and blameworthiness. If they fail to act as we would like them to, they do no wrong: they merely malfunction.

Autonomous vehicles may face some of the hypothetical dilemmas that appear in non-consequentialist theorizing. For example, consider this variant on *Ambulance*:

*Ambulance*\*\*

A self-driving ambulance carrying five injured persons finds that it is unable to get to the hospital in time to save any of the five unless it runs over a person who is in its path.

We may want autonomous vehicles to come programmed with instructions that cover dilemmatic cases such as these. In particular, it may seem plausible that we should want those instructions to accord with what a human being morally ought to choose if she were at the wheel.<sup>18</sup> For example, in *Ambulance*\*\*, we may want it to be the case that the ambulance allows the five to die rather than proceed ahead and kill the one.

Suppose, however, that we accept the *Wrong-Preventing Principle* as central to the explanation of why we should prevent optimistic wrongs. Then we do not have decisive reason to ensure that autonomous cars behave in this way. If we accept the *Wrong-Preventing Principle* as crucial to explaining the difference between *Footbridge*\* and *Tracks*, we

<sup>17</sup> See Wendell Wallach and Colin Allen, *Moral Machines: Teaching Robots Right and Wrong* (Oxford, 2009). More generally, as robots and drones become more prevalent they will increasingly be required to respond to morally relevant information about their surroundings. Especial concern surrounds the deployment of robots in war, on which see Ronald Arkin, *Governing Lethal Behavior in Autonomous Robots* (Boca Raton, 2009).

<sup>18</sup> With respect to the morality of deploying autonomous robots in battle, it seems to be taken as given that they should conform to the laws of war that govern human beings. See e.g. Arkin, *Governing Lethal Behavior in Autonomous Robots*. These laws arguably reflect non-consequentialist principles, such as the intention/foresight distinction: see McMahan, ‘Intention, Permissibility, Terrorism, and War’.

suppose that our obligations to prevent optimific wrongs hinge on the intrinsic desirability of preventing wrongful acts. Thus, since impeding an autonomous car would achieve no such end, there is no obligation to do so. *Ambulance\*\** would be like *Tracks*: the needs of the many should outweigh the needs of the few.

Many people will no doubt be unhappy with the conclusion that there is no decisive reason to prevent self-driving cars from killing people in situations like *Ambulance\*\**. What we think about cases like this may be thought to determine how we regulate the autonomous vehicles of the near future. We may feel uncomfortable at the suggestion that we have no reason to recall and rewire a fleet of autonomous cars that are discovered to employ a purely consequentialist 'moral compass'.

To make matters worse, on this point the *Wrong-Preventing Principle* asks us to draw distinctions that strike us as peculiar. If one came across the situation described in *Ambulance\*\** and saw the car proceeding along, the principle suggests that one should have to check first whether the car is self-driving or person-controlled. After all, if a person is at the wheel, then one ought to intervene; but if the car is autonomous, intervening would be like saving the hiker and killing the five in *Tracks*. It would seem to be of considerable moral importance that one be able to ascertain correctly the nature of the vehicle's operator. One might be required to take considerable care in this matter before deciding whether to intervene. The idea that we should have to decide this issue before deciding to intervene in a case like *Ambulance\*\** is one that I expect most people will find quite puzzling.

Here is the third argument against supposing that the *Wrong-Preventing Principle* can explain the difference between *Footbridge\** and *Tracks*. As noted, it is similar in character to the second argument, and so my presentation will be much shorter. The *Wrong-Preventing Principle* gives us no reason to prevent optimific wrongs that prevent other wrongs of equal or greater magnitude. For example, if we modify *Footbridge\** so that the trolley is not out of control but within the control of someone who means to murder the five, then our duty to intervene would evaporate. In parallel with my second argument, I believe that most people who are inclined to intervene in *Footbridge\** will find this implication counterintuitive. Moreover, I expect that they will find it intuitively surprising that this difference should make such a difference. It seemed an incidental detail of the set-up in *Footbridge\** that the trolley was said to be out of control rather than directed by a person.

This completes my case against the attempt to deny premise 2 in my argument by appeal to the *Wrong-Preventing Principle*. As I hope to have shown, this strategy is implausible in a number of respects. First, the *Wrong-Preventing Principle* can lead us to fetishize the prevention

of injustices over the welfare of human beings. Second, it implies that we may allow acts that are in all respects like optimific wrongs except that they involve agents who lack the moral capacities required for blameworthiness; our intuitions appear to recognize no distinction here. Third, it implies that we may allow optimific wrongs if they prevent other wrongs; once again, it appears that our intuitions fail to recognize such a distinction. All in all, it seems we just do not value the prevention of wrongful acts in the manner that the appeal to the *Wrong-Preventing Principle* suggests. Therefore, we cannot appeal to the *Wrong-Preventing Principle* to resist premise 2.

We appear, then, to have no means left by which to resist my argument: each premise seems true and the conclusion follows straightforwardly. We thus have reason to believe that when it comes to optimific wrongs, the distinction between action and omission may make the same kind of difference as obtains between *Ambulance* and *Ambulance\**.

#### IV

I have argued against the view that we ought to prevent optimific wrongs. The key premise in my argument was the assertion of parity between *Footbridge\** and *Tracks*, which I defended by dismissing the only minimally plausible moral principles that might be used to drive a wedge between the two cases. In section II, I set aside the *Scope Principle* rather quickly, as this principle seemed to draw any support it might have from a common cognitive bias. In section III, I then set out three arguments that challenge the suggestion that we should appeal to the *Wrong-Preventing Principle* to differentiate between *Footbridge\** and *Tracks*. The cumulative effect of these arguments was to cast significant doubt on the suggestion that we attach this kind of importance to the prevention of wrong acts. Having set aside both the *Scope Principle* and the *Wrong-Preventing Principle*, I believe we should conclude that my argument in section II is sound.<sup>19</sup>

[andreas.mogensen@philosophy.ox.ac.uk](mailto:andreas.mogensen@philosophy.ox.ac.uk)

<sup>19</sup> For helpful comments on previous drafts of this article I am grateful to Krister Bykvist, William MacAskill, and the audience at the Balliol Positive Ethics Seminar.