

LESSER-EVIL JUSTIFICATIONS FOR HARMING: WHY WE'RE REQUIRED TO TURN THE TROLLEY

BY HELEN FROWE

Much philosophical attention has been paid to the question of whether, and why, one may divert a runaway trolley away from where it will kill five people to where it will kill one. But little attention has been paid to whether the reasons that ground a permission to divert thereby ground a duty to divert. This paper defends the Requirement Thesis, which holds that one is, ordinarily, required to act on lesser-evil justifications for harming for the sake of others. Cases in which we have lesser-evil justifications of harming for the sake of others are rescue cases. Ordinarily, an agent is under a duty to rescue unless doing so imposes too great a cost on her, or violates someone else's rights. When neither of these defeating conditions obtain, one is required to rescue even if this involves causing harm to innocent people.

Keywords: lesser-evil justifications, harming, self-defence, trolley problem, bystanders, duties to rescue.

I. INTRODUCTION

Lesser-evil justifications invoke the disparity between the harm that one will inflict on an agent, and the harm that one will prevent by harming her.¹ On a non-consequentialist framework, when this disparity is sufficient to overcome the moral worseness of causing harm compared to allowing harm, causing harm is justified as the lesser evil (Feinberg 1972; Thomson 1986). Unlike other justifications for harming, such as liability, consent or desert, lesser-evil justifications do not rely upon the idea that the agent being harmed lacks a right not to be harmed. They rely solely on the amount of harm one can prevent.

Consider *Trolley*:

¹ In line with much of the literature, I'll set aside here cases in which one might produce benefits by harming a person. I take it that our rights to be benefitted are much less stringent than our rights not to be harmed. The thesis defended here might also hold in these cases, but I do not explore this possibility here.

Trolley: A runaway trolley is heading to where it will kill five people. Pedestrian is standing next to a switch that will divert the trolley down a sidetrack, saving the five. However, the trolley will then kill Workman, who is trapped on the sidetrack.²

Most people endorse the following three claims about *Trolley*. The first is that Pedestrian has a lesser-evil justification for killing Workman (Hallborg 1997; Kamm 2007; McMahan 2014). The second is that, if she flips the switch, Pedestrian will permissibly infringe, rather than violate, Workman's right not to be killed (Rodin 2017: 30; Thomson 1986).³ (One *infringes* a right when one transgresses that right with justification; one *violates* a right when one transgresses that right without justification.) The third claim is that Pedestrian does not *have* to flip the switch. She may refrain from saving the five (Alexander 2005: 616, 619; Walen and Wasserman 2012). This paper argues that this third claim—that Pedestrian may decide whether or not to act on the lesser-evil justification for harming Workman—is false.

Lesser-evil justifications for harming are an important part of non-consequentialist moral theories (Bazargan 2016; Frowe 2014; McMahan 2009; Quong 2015; Rodin 2011; Tadros 2011). Any form of threshold (that is, non-absolutist) deontology is likely to at least sometimes rely on them to explain why one may do harm despite the moral worseness of harming compared to allowing harm. They frequently crop up in political philosophy and in the law (for example, as 'necessity' justifications for putative wrongdoing). That we have a criminal legal system at all, despite knowing that we will sometimes wrongly convict the innocent, is typically justified on lesser-evil grounds. These justifications also play an important role in the ethics of war. They are often thought to explain the permissibility of inflicting collateral harms on civilians (Lazar 2014: 12; Moore 2014: 198; Walzer 1977: 153). They may also do significant work in explaining the permissibility of intentionally killing combatants in war (Lazar 2015: 10; McMahan 2009: 45). More generally, war itself is plausibly justified only if it is the lesser evil compared to failing to fight (McMahan 2009: 21, 45).⁴ And yet, despite their importance, our treatments of lesser-evil justifications for harming—especially compared to the various book-length treatments of liability, consent and desert-based justifications for

² The original trolley case was proposed by Foot (1967). The variation of the bystander at the switch comes from Thomson (1985: 1397). In all the ensuing cases, I assume that nobody involved bears any responsibility for the threat that facing the five, nor has any special relationship to another party.

³ That Workman's right is infringed does not follow analytically from the fact that one has a lesser-evil justification, one might think that, for example, agents simply lack rights against being harmed when harming them is justified as the lesser evil. I take no stance on this here.

⁴ Oberman (2015) has recently argued that there are no morally optional wars, but not the grounds that we are generally required to act on lesser-evil justifications, as argued here. Oberman's arguments are distinctively political, and explicitly restricted to the behaviour of states.

harming—remain distinctly hand-wavy. This paper aims to say something more concrete about these justifications.

Any plausible account of lesser-evil justifications is likely to hold that one is *sometimes* required to act on such justifications: one would, for example, be required to painfully scratch an innocent person to save a million lives. Likewise, there are some cases in which one is plausibly permitted but not required to act to secure the lesser evil. For example, one need not act on a lesser-evil justification if so acting imposes a considerable cost on oneself—for example, if turning the trolley would involve sacrificing one's own life. Nor need one act on justifications for harming for one's own sake, since the right of self-defence is discretionary.

One might also face what we can call a *parity case*, in which the amount of harm one can prevent is, morally speaking, exactly on a par with the harm one will cause (taking into account that, by acting, one would be causing rather than allowing harm). Here too, one would be permitted to either act or refrain from acting.⁵ (I include parity cases as lesser-evil cases because they share the basic structure of lesser-evil justifications—that is, the justification for harming is grounded solely in the amount of harm that one can thereby prevent. Alternatively, we might restrict the term 'lesser evil' to cases in which the harm prevented *outweighs* the harm inflicted. Here, I'll treat parity cases as an anomalous type of lesser-evil case that I will exclude from the scope of my argument. But we could equally think that they fall outside its scope because they are not lesser-evil cases at all).

Although I will assume that *Trolley* is not a parity case, anyone who thinks otherwise should simply increase the number of lives to be saved from five to six. If the moral reasons against harming Workman are exactly balanced by the moral reasons to save five, such that morality is indifferent between harming and not harming, the balance must tip in favour of harming Workman when one will thereby save six.

I set aside here these largely uncontroversial cases. The interesting disagreement arises in cases such as *Trolley*, where (a) the agent acts for the sake of others, (b) there is sufficient, but not overwhelming, disparity between the relevant harms, and (c) there is no physical or material cost to the intervening agent of the sort that might render acting supererogatory. Call these *ordinary cases*. As noted above, the standard view is that acting on a lesser-evil justification is morally optional in ordinary cases—that, for example, Pedestrian may either kill or refrain from killing Workman. This optionality is usually explained by invoking the moral significance of the (potentially) intervening agent's autonomy, the protecting of which gives her a right to refrain from harming.

⁵ Thanks to Ralph Wedgewood for helpful discussion of this point.

I will call this view—that, in ordinary cases, bringing about lesser-evil harms is merely permissible—the *Permissibility Thesis*.

The Permissibility Thesis: in ordinary cases, an agent's autonomy interests render acting on lesser-evil justifications for harming for the sake of others morally optional, even though there is no risk of unreasonable physical harm to her.

I reject the Permissibility Thesis, arguing that an intervening agent's autonomy interests do not render acting on lesser-evil justifications morally optional. I defend the Requirement Thesis:

The Requirement Thesis: In ordinary cases, an agent is morally required to act on lesser-evil justifications for harming for the sake of others, unless (a) acting would impose some very significant psychological cost on her, or (b) an alternative action would be equally or less harmful overall.

Given the pervasiveness of lesser-evil justifications in our moral, political and legal systems, whether we ought to adopt the Requirement Thesis or Permissibility Thesis is of considerable importance.⁶

The Requirement Thesis is primarily motivated by the recognition that cases in which we have lesser-evil justifications for harming are rescue cases, and should be treated as such. Ordinarily, we have a duty—and not a mere permission—to rescue, unless doing so imposes too great a cost on us, or violates someone else's rights. This duty correlates with a claim to be rescued on the part of the prospective victim(s). The Requirement Thesis is thus underpinned by the following more general principle, *Preventing Harm*:

Preventing Harm: One has a duty to prevent harm to others when one can do so without violating anyone's rights, and without bearing an unreasonable cost.

I do not try to establish here an exact threshold for an unreasonable cost. Rather, I seek to establish that autonomy interests cannot ground an exception to the duty to rescue. Thus, we should reject the Permissibility Thesis in favour of the Requirement Thesis.

The paper proceeds as follows. Section II elaborates on the nature of lesser-evil justifications. Section III motivates the Requirement Thesis by contextualizing lesser-evil cases as rescue cases. It further argues that the Requirement Thesis holds even when one bears a supererogatory cost to be in a position to rescue. It also defends *Preventing Harm* against possible objections. Section IV rejects the idea that one has sufficient reason to either rescue or refrain from rescuing in ordinary lesser-evil cases. It also explains the caveat that one

⁶ Thomson (1985) defends the claim that the bystander is permitted to turn the trolley, but not required to do so. In Thomson (2008), she retracts her claim that turning is permissible for the bystander (although suggests it is required for the driver). FitzPatrick floats the idea that diverting is required, but does not defend it (FitzPatrick 2009: 636). Alexander argues that diverting is merely permissible unless one has an agent-relative duty (Alexander 2005: 618).

is required to secure the lesser evil only if there is no equally or less harmful alternative. Section V argues that autonomy interests in shaping our own lives do not undermine the Requirement Thesis. Section VI considers the implications of the Requirement Thesis for the defensive rights of those who are harmed on lesser-evil grounds. Section VII concludes.

II. LESSER-EVIL JUSTIFICATIONS FOR HARMING

Although lesser-evil justifications depend to some extent on the consequences of acting or refraining from acting, they are not (as I will understand them here) purely consequentialist justifications.⁷ They are also sensitive to deontological constraints on action, such as the moral significance of the distinction between doing and allowing. One must prevent significantly more harm than one causes in order to overcome the presumptive worseness of causing harm compared to allowing harm (e.g. McMahan 2009: 28). Thus, most non-consequentialists believe that one may not, for example, divert a trolley away from two people to where it will kill one, even though this would result in the fewest deaths. Saving two lives is not a sufficient good to defeat the presumptive worseness of killing compared to letting die.

Lesser-evil justifications are also sensitive to modes of agency—for example, to whether a person is to be harmed as a foreseen side-effect, or opportunistically (that is, in a way that enables one to benefit from her presence, such as by using her as a human shield). It is easier to justify harming a person as a foreseen side-effect than it is to justify opportunistically harming her (Frowe 2008; Lazar 2015: 56–73; Tadros 2011: 113–16; Thomson 1985). For example, most non-consequentialists believe that one may not push someone in front of a trolley in order to save five lives, even if one may divert the trolley away from five to where it kills one as a side-effect. Preventing five deaths is insufficient to justify opportunistically killing one.

Noticing the range of features to which lesser-evil justifications are sensitive is important because it forestalls the objection that when we talk about ‘the lesser evil’, we are merely comparing the outcomes of our available courses of action. This objection holds that we might grant that one death is of course better than five deaths—it is the lesser evil—but that this observation does not by itself justify any particular course of action, since there are deontological constraints on actions that seek to minimize harm. But this objection is misguided. As David Rodin argues, ‘the value of states of affairs relevant to lesser evil justification depends crucially on how harms and benefits come into being, and how they are related to human agency’ (Rodin 2011: 97; see also

⁷ Although I will assume a non-consequentialist framework, much of what I say could likely be adopted by at least some consequentialists.

Alexander 2005: 616). To say that an action is justified as the lesser evil is to take into account not only its outcome, but also its deontic features.

III. LESSER-EVIL CASES AS RESCUE CASES

III.1. *Duties to rescue*

The standard view of cases such as *Trolley* is that harming Workman is objectively justified—that is, the balance of moral reasons favours killing him so that the five are saved (recall our assumption that *Trolley* is not a parity case) (McMahan 2014: 135; Quong 2016: 826, n.23). And yet the Permissibility Thesis holds that Pedestrian may refrain from engaging in this justified rescue if she prefers that the five die rather than that she kill the one.

The Requirement Thesis, in contrast, holds that Pedestrian must divert the trolley. The underlying rationale for this claim is simple—namely, that Pedestrian finds herself in a situation in which she can save several people's lives without violating anybody's rights and without bearing unreasonable costs. I propose the following principle:

Preventing Harm: One has a duty to prevent harm to others when one can do so without violating anyone's rights, and without making a supererogatory sacrifice.

The duty to prevent harm in *Trolley* is owed to the five, who have a corresponding right to be saved when saving them does not require the violating of anyone's rights, nor involve anyone's making a supererogatory sacrifice.⁸

This duty obtains even when saving involves infringing rights as the lesser evil. Consider *Storm*:

Storm: A hiker collapses in a storm. A second hiker can pull her inside a cabin, but only by stepping on the cabin owner, who is sleeping deeply in front of the door, safe in his thermal sleeping bag. Stepping on the owner will break the owner's hand.⁹

The second hiker may not leave the first to die on the grounds that she would prefer not to infringe the cabin owner's rights. Rather, she must step on the owner and pull the unconscious hiker inside. Doing so imposes no great burden on her and violates nobody's rights, since the cost to the cabin owner

⁸ Quong suggests that granting the five a right to be saved relies on a kind of aggregation that is at odds with how we think of rights (Workshop on Self-Defence, USC, April 2015). But individuals can acquire rights on the basis of aggregating interpersonal interests, since aggregating interests can render otherwise disproportionate defence proportionate (see Frowe 2014, 2015). The permission to kill Workman is also a type of right. If interests can aggregate to generate permissions, they can plausibly aggregate to generate claims.

⁹ Adapted from a case in Feinberg (1972: 102). Feinberg doesn't use the language of lesser evil: I merely borrow his case as an example of when one would be required to infringe rights on lesser-evil grounds.

is justified as the lesser evil. In such circumstances, there is a duty to save, not a mere permission.

III.2. *Supererogatory rescues*

I've restricted the Requirement Thesis to ordinary cases, in which acting on a lesser-evil justification does not require the rescuer to bear an unreasonable physical cost. One cannot, of course, be required to bear an unreasonable—that is, supererogatory—cost. But we should note that the Requirement Thesis nonetheless applies in cases where one elects to bear a supererogatory cost and thereby finds oneself in a position to enact a lesser-evil rescue.

The view that one can incur conditional duties if one chooses to bear a supererogatory cost is familiar. One need not risk one's life by rushing into a burning building. But, if one does rush in, one may not save the caged bird instead of the child (Kagan 1989: 16; Parfit 2011: 225). The Effective Altruism campaign is premised on the view that even supererogatory donations to charity should be directed to where they will minimize harm (MacAskill 2015; Pummer 2017). Similarly, the Requirement Thesis holds in cases such as *Flower Trolley*:

Flower Trolley: A trolley is heading towards five people. There is a switch that will divert the trolley down a side-track, on which Workman is trapped. To get to the switch, Pedestrian must climb an electric fence that will cause her to be paralysed from the waist down. Pedestrian climbs the fence so that she can take a closer look at a rare flower. Once she drags herself to the flower, she can pull the switch.

Pedestrian is not required to paralyse herself to save the five. But if she decides to climb the fence, she must now save the five. That she has borne a supererogatory cost to be in a position to save does not alter her duties with respect to saving. Just as supererogatory rescuers are still under a duty to minimize harm when saving (as the Effective Altruists claim), Pedestrian is still under a duty to infringe rights as the lesser evil, as the Requirement Thesis claims.

III.3. *Infringing rights*

I take it that proponents of the Permissibility Thesis are not generally sceptical of duties to rescue—at least, they need not be, and I will not defend the existence of such duties here.¹⁰ But they might nevertheless object that *Preventing Harm* begs the question against their view, since it holds that saving can be required even if it involves infringing rights. They might propose an alternative principle:

¹⁰ Even the more libertarian-inclined Permissibility Theorists, such as Alec Walen, grant that we have some duties to aid. See, e.g. Walen (2014: 454).

*Preventing Harm**: One has a duty to save only when one can prevent harm to others without violating or infringing anyone's rights, and without making a supererogatory sacrifice.

But few people subscribe to the view that one is *never* required to infringe rights in the course of saving others. It seems clearly impermissible to refuse to save the stranded hiker, for example, or to scratch one to save a million. And even though killing is harder to justify than those harms, any plausible view will grant that there is *some* amount of preventable harm that could make killing required. If so, we cannot point to the fact that *Trolley* involves infringing rights to show that saving is not required, and we cannot endorse *Preventing Harm**.

What if we insist that one cannot be required to infringe rights—or, at least, to infringe the right not to be killed? Perhaps, contrary to most people's intuitions, we ought not to divert the trolley away from a thousand if it will then kill one. This reply demands an account of the morally significant features of killing that mark it out as the kind of action that can never be required as the lesser evil. But it is hard to see how such an account could be consistent with the permissibility of diverting the trolley. Killing is surely no less objectionable in virtue of the fact that a person volunteers to kill rather than finding herself under a duty to kill. The upshot of such an argument would be that killing is simply never the lesser evil, not that killing is discretionary. It would be a rejection of the Permissibility Thesis, not a vindication of that thesis.

Proponents of the Permissibility Thesis thus face a dilemma. They can grant that infringing rights, including the right not to be killed, is sometimes required as the lesser evil. This leaves *Preventing Harm* intact. Or, they can deny that killings that infringe rights are ever required because of the character of killing, thereby undermining the Permissibility Thesis itself. Since the latter is clearly unappealing, I think they must grant *Preventing Harm*.

IV. MORAL PARITY AND SUFFICIENT JUSTIFICATION

One might object to the Requirement Thesis by invoking a particular account of moral permissibility, according to which an action is morally permissible if there is sufficient justification for performing it, which could be true of a range of mutually exclusive actions.¹¹ In *Trolley*, we agree that there is sufficient reason to divert, since diverting saves five. But perhaps there is also sufficient reason to refrain from diverting, because diverting kills one. This makes either diverting or refraining from diverting permissible.

However, this is true of *Trolley* if, and only if, it is a moral parity case of the sort described in Section I. In parity cases, there is sufficient justification for diverting or refraining because there is equal reason to divert or refrain.

¹¹ I owe this suggestion to an anonymous referee.

But once we move beyond a case of parity—say by increasing the preventable harm—there must be *more* reason to harm than to refrain from harming. When one has a range of available actions, and most moral reason to perform a particular action, one cannot insist that, nonetheless, the alternatives are justified. Justification is comparative: as Pummer argues, the moral status of an action is sensitive to the available alternatives (Pummer 2017: 87). For example, consider *Two Trolleys*:

Two Trolleys: Trolley A is heading towards five people. Pedestrian can divert it to where it will kill Workman. Trolley B is heading towards ten people. Pedestrian can divert it to where it will kill Engineer. She does not have time to divert both trolleys.

The sufficient justification account implies that Pedestrian may divert Trolley A, saving five lives, instead of diverting Trolley B, saving ten lives. But this seems clearly impermissible. That, other things being equal, saving five lives suffices to justify killing Workman does not show that killing Workman is always justified. Specifically, it is not justified given the alternative of saving ten and killing Engineer—that is, of doing more good.

Cases such as *Two Trolleys* explain the Requirement Thesis's caveat that one must act on lesser-evil justifications only if no alternative action would be equally or less harmful overall. This might seem like an odd caveat: we might think that acting can be the lesser evil only if there's no less harmful alternative. But *Two Trolleys* suggests otherwise. There is, in a sense, a lesser-evil justification for killing Workman. Killing him saves five lives, as in *Trolley*. But there isn't a lesser-evil justification for diverting Trolley A rather than Trolley B. On the contrary, diverting Trolley A conflicts with the duty to minimize harm when saving. In my view, the duty to rescue is owed to those the saving of whom will minimise harm—that is, Pedestrian has a duty that she owes to the ten to save them. Failing to divert Trolley B will therefore violate the ten's rights to be saved. It is thus prohibited by *Preventing Harm*.

Of course, the foregoing doesn't mean that one must always do what one has most moral reason to do: some actions are so costly as to be morally optional. But, as I'll now argue, acting on lesser-evil justifications for harming is not typically optional. More specifically, so acting is not rendered optional by the compromising of a rescuer's autonomy interests.

V. AUTONOMY, INTERESTS AND THE DUTY TO RESCUE

The Permissibility Thesis holds that a rescuer's autonomy underpins the optionality of killing Workman. For example, Alec Walen and David Wasserman suggest that the key to resolving trolley cases lies in recognizing Pedestrian's claims, as the rescuer, along with the competing claims of Workman and the five, as potential victims or beneficiaries of Pedestrian's actions. They grant

that if only the claims of the five and the claims of Workman were relevant, then Pedestrian would indeed be required to divert the trolley. But, they argue, this approach mistakenly treats Pedestrian as a mere ‘node’ for mediating the claims of others, neglecting the fact that Pedestrian also has claims—most importantly, a claim not to act in ways in which she does not want to act. This claim protects her from having to treat herself as a tool for the promotion of the good. Since diverting the trolley would kill Workman, and Pedestrian ‘has a substantial interest in not having to be an agent of death’, her claim is powerful enough to make diverting discretionary (Walen and Wasserman 2012: 554). Thus there is no requirement to act on lesser-evil justifications in the ordinary cases with which we are concerned.

We can helpfully distinguish two components of the autonomy cost that Walen and Wasserman identify. The first is the cost to the agent’s interest not being an agent of death (I’ll focus on death but, *mutatis mutandis*, the following also applies to other harms). The second is the cost of being required to treat oneself as a tool for the good of others. I’ll tackle each of these in turn.

V.1. On (not) being an agent of death

There are two ways in which we might construe being an agent of death as compromising an agent’s interests. First, if not being an agent of death is an objective, preference-independent interest, then making someone kill will inevitably compromise that interest. Secondly, even if not being an agent of death is not an objective interest, it might still compromise an agent’s autonomy interest—that is, the interest she has in shaping her life and pursuing her freely chosen ends. The ends themselves need not be intrinsically valuable, but the fact that she has chosen to pursue them is valuable. So, if an agent has adopted an end of not killing people, forcing her to kill someone will compromise her autonomy interests.

Interestingly, Walen and Wasserman’s account straddles these two types of interest. On the one hand, they imply that the interest in not being an agent of death is an objective interest, insofar as they assert without argument that we all have such an interest, and they take the Permissibility Thesis to apply to agents quite generally (2012: 554). But, on the other hand, compromising this interest wouldn’t obviously compromise autonomy, despite their focus on an agent’s claim not to act in ways she does not want to act. However, as I’ll argue, neither construal of the interest in not being an agent of death vindicates the Permissibility Thesis, so this need not detain us.

There are, plausibly, some interests that we have independently of our preferences, such as interests in avoiding physical harm. People disagree about the range of these objective interests. For example, some people think that it is objectively good for me to read certain books, watch certain plays and view

certain works of art. Others might restrict the range of objective interests to the things I need in order to live a minimally decent life.

If the interest in not being an agent of death is an objective interest, it is probably best characterized as a moral interest. Plausibly, we all have an interest in living a morally good life—other things being equal, a life in which I do not rape or steal and so on is objectively better for me than a life in which I do those things.

But I doubt that there is a moral interest in not being an agent of death. Killing is sometimes permissible and, plausibly, sometimes morally required. I may kill someone in self-defence, or euthanize a terminally ill patient. Someone who refuses to lethally defend a child—especially their own child—against a culpable murderous attacker, or to divert a trolley from a thousand people to one, does something seriously morally wrong. While we clearly have a moral interest in refraining from *unjustified* killing, we have no similar interest in refraining from justified killing. The same is true of aversions to killing. An aversion to unjust killing is surely something we should all cultivate. But it doesn't follow that we should extend this aversion to killing generally.

Even if we grant that we each have a moral interest in refraining from killing generally, we must also grant that we each have a moral interest in saving lives. Of course, many people think that this interest is less weighty than our interest in refraining from killing. But the strength of any moral interest in avoiding killing would surely track the moral disvalue of killing, and the strength of any moral interest in saving lives would similarly track the moral value of saving lives. If so, the balance of these interests in *Trolley* would favour diverting even granting the greater importance of avoiding killing, since (by hypothesis) the value of saving the five outweighs the disvalue of killing Workman. Walen and Wasserman agree that, setting aside Pedestrian's interests, the moral reasons against killing Workman are outweighed by the moral reasons for saving the five. It would be strange to nonetheless believe that Pedestrian's moral interest in saving the five is outweighed by her moral interest in not killing Workman. The greater compromise of Pedestrian's moral interests—understood as an interest in acting in accordance with the balance of moral reasons—would be failing to save the five.

Notice too that if we deny that this is how the reasons balance out, and insist that Pedestrian's moral interest in not being an agent of death somehow outweighs her moral interest in saving the five, we are again at a loss to explain even the permissibility of diverting. One may not decide to go on a stealing spree because one has decided to waive one's interest in not being the sort of person who steals things. If the balance of moral reasons comes down in favour of not stealing, one simply ought not to steal. Similarly, if the interest in not being an agent of death is grounded in moral reasons not to kill, one cannot simply decide to discount these reasons in one's own case. We should not characterize Pedestrian as having a claim not to kill on this account, but

rather as having a duty not to kill. One cannot simply waive one's moral duties. This argument for the Permissibility Thesis thus looks unpromising.

V.2. *Autonomy, and treating oneself as a tool*

Perhaps autonomy interests will be more fruitful as a defence of the Permissibility Thesis. We might think that Pedestrian is entitled to decide whether her life is one that involves killing—or, more specifically, whether her life is one that involves killing Workman. According to this view, it is the moral value of letting Pedestrian choose how to shape her life—rather than the disvalue of killing Workman—that grounds her permission to refrain from diverting the trolley.

But, of course, we may not shape our lives however we please. The scope of the ends that I am allowed to adopt for myself, and when and how I may pursue those ends, is constrained by my moral duties. Faced with a drowning child, the fact that I have chosen to shape my life in a way that does not involve rescuing children would not make it permissible for me to fail to save her. Likewise, that Pedestrian has chosen to shape her life in a way that does not involve killing Workman does not show that she may refrain from killing him.

The moral significance of an agent's autonomy interests—and the stringency of the obligation to let her pursue those interests—plausibly tracks the extent to which compromising such interests will affect her capacity to shape her life. If requiring her to act will significantly compromise that capacity, then we may impose such a requirement only if the moral good at stake is very important indeed. But if requiring her to act will have little impact on her capacity to shape her life, we can impose such a requirement for less important goods.

Agents are clearly *sometimes* required to compromise their autonomy interests for the benefit of others. The second hiker must drag the first into the cabin whether she wants to or not. If this is what it means to make a person treat herself as a tool, it follows that making people treat themselves as tools is sometimes permissible. This includes requiring agents to cause harm to innocent people even if they would prefer not to. If diverting the trolley will break Workman's leg, rather than kill him, Pedestrian surely cannot invoke her autonomy interests to justify not diverting the trolley. Even if she has adopted an end of not causing harm to people, and tries to shape her life accordingly, the compromising of her interest in pursuing her ends is outweighed by the harm she can prevent to the five.

Given this, the issue at hand is whether requiring someone to kill an innocent person is a sufficiently serious impeding of her capacity to shape her life that we may only rarely impose such a requirement on her, when doing so averts an utter catastrophe. Obviously, it is trivially true that if a person has adopted an end of living a life in which she does not kill anyone, making her kill someone

will make it impossible for her to shape her life in that way. But the same would be true of someone who has adopted an end of not rescuing drowning children, and I doubt we would be much swayed by the argument in that case. The question is not whether a requirement to act in a given situation makes it impossible for the agent to pursue any particular end, but whether it restricts her ability to shape her life in a more general way.

Compare *Trolley* with *Enlisting*:

Enlisting: Sarah can enlist in the military, where she will, over the course of a ten-year career, contribute to the collateral killing of ten innocent people. However, she will also contribute to the saving of fifty innocent lives. But Sarah does not want to join the military. She wants to be an artist (even though she will not be a very good one).

The ratio between the number of lives saved and the number of people killed in *Enlisting* is the same as in *Trolley*. The difference is the cost to Sarah compared to the cost to Pedestrian. Requiring Sarah to join the military would very significantly restrict her ability to pursue her own ends not only for the substantial period of her life that she would spend in the military, but also afterwards, when she will have sacrificed ten years of her development as an artist. The objection to a requirement to enlist, then, is not that we will be requiring Sarah to be an agent of death, but that by making her enlist, we drastically restrict the sphere in which she gets to shape her own life.

It seems plausible that Sarah may refuse to join the military. Enlisting is so costly that doing so is supererogatory. But it is not plausible that killing as a justified side-effect will generally restrict an agent's capacity to shape her life to the extent that it renders killing supererogatory. Some killings are courageous or heroic; others are mundane or institutionalized. Someone might have a career as police sniper, for example, without that impeding her ability to live a good life. Acts of killing, and people's responses to them, are far too varied to support the view that justified rescue killings in ordinary cases are always so costly to autonomy as to be supererogatory.

This is, in part, because the impact of killing on a person's autonomy would have to be very significant to defeat the duty to save five lives. We can see this by comparing the harms at stake in *Trolley*:

- (1) Killing Workman as a side-effect of an action that benefits others
- (2) Allowing each of the five to die
- (3) Requiring Pedestrian to treat herself as a means, killing someone else as a side-effect of saving five

The harm to Pedestrian's autonomy is very unlikely to be worse than the lethal harms facing Workman and the five, even taking into account that Pedestrian is harming herself as a means. Being required to kill or seriously harm oneself as a means is worse than being killed as side-effect. But I doubt that being required to treat oneself as a means by seriously harming or killing *someone else*

as a side-effect is worse than being killed as a side-effect oneself, which is what we are assuming it is permissible to do to Workman.

Bear in mind, too, that it would be permissible for the five to divert the trolley towards Workman themselves. I suggest that we cannot plausibly grant the five rights of defence even though they will kill Workman, but deny them a right to be saved because of the harm to Pedestrian's autonomy. This seems objectionably dismissive of the harm we force Workman to bear: it allows the five to proceed in the face of Workman's objection that they will kill him, but insists that Pedestrian's protection of her autonomy defeats their claims. If the five may proceed despite the lethal harm to Workman, Pedestrian must proceed despite even somewhat serious harm to her autonomy.

Only if Pedestrian will suffer a very serious setback of her autonomy—of the sort that will significantly affect her ability to shape the rest of her life in accordance with her freely chosen ends—could requiring her to impose this harm on herself render the rescue of the five supererogatory.¹²

This is not to say that only costs that significantly compromise one's ability to live a flourishing life can defeat a requirement to act on lesser-evil justifications. Physical costs—such as the loss of a hand—can render a rescue supererogatory even if one could nevertheless go on to have a good life overall. This is because our rights against physical harm are primarily grounded in our interests in bodily integrity, and don't (solely) depend on whether the harm will affect one's overall flourishing. It would, for example, be very wrong to gratuitously shove a needle under your fingernail for a minute, inflicting horrible pain on you, even if you would go on to have a perfectly normal life afterwards.

Since psychological harms do not threaten our interests in bodily integrity, the threshold at which psychological harms become supererogatory is not when they compromise flourishing to the same degree as a supererogatory physical harms. Thus, Pedestrian is not exempted from the duty to rescue if, for example, rescuing compromises her flourishing to the same degree as the loss of a hand. Rather, without the further threat to bodily integrity, only very serious compromising of flourishing defeats a duty to rescue.

To get a sense of how serious the harm to Pedestrian's autonomy would have to be, consider how much prospective harm to autonomy it would take to defeat the duty to save a drowning child. After all, autonomy costs can arise irrespective of whether or not a rescue involves infringing rights: what matters is that Pedestrian is being made to do something she does not want to do. Pedestrian may be so concerned about overpopulation, for example, that she will find it genuinely traumatic to save a life, having refrained from having children herself and spent her life campaigning against reproduction in general. But, I suspect, this trauma would have to pose a long-term and

¹² See, Fabre (2007: 370) for a discussion of psychological costs and the duty to rescue. I borrow the term 'psychological' from Fabre's characterization of this kind of cost.

life-altering threat to her autonomy before we might even begin to think it grounds for failing to save the child.

This is one way in which autonomy-based defences of the Permissibility Thesis prove too much. Presumably, proponents of the Permissibility Thesis do not believe that one may ordinarily refrain from saving a child even if one is very concerned about overpopulation. They must be cautious, then, about suggesting that one may refrain from saving the five if one is very keen to be a pacifist. One could have autonomy-based interests in both these things, but those interests are generally unlikely to defeat duties to rescue. This worry extends to duties to rescue innocent people from liable attackers by engaging in other-defence.¹³ That the attacker's right is not infringed, since she lacks a right not to be harmed, might make no difference to how the defender feels about killing her. Relying on the effect of killing on the agent risks labelling all killings, and some rescues that don't even involve killing, supererogatory.

Thus far, I have argued that by approaching lesser-evil cases as rescue cases, we find strong support for the Requirement Thesis. It is common ground between the Permissibility Thesis and the Requirement Thesis that one is permitted to save the five: ours is not a debate about whether killing Workman is justified. Given the agreed permission to rescue, it is then natural to move to a presumption of a duty, as normally obtains in rescue cases. It is now the failure to rescue that must be justified.

I have defended *Preventing Harm*, which holds that one has a duty to prevent harm to others when one can do so without violating anyone's rights, and without making a supererogatory sacrifice. I suggested that proponents of the Permissibility Thesis cannot plausibly deny that one is sometimes required to infringe the right not to be killed, and so they should not reject *Preventing Harm* on these grounds. Moreover, any insistence that one is never required to infringe the right not to be killed would have to be underpinned by an account of the features of killing that explain why it can never be so required. But, I argued, this account will undermine the permissibility of diverting the trolley.

I suggested that the most promising avenue for defending the optionality of lesser-evil cases is to show that the rescuer's autonomy interests protect her from having to enact rescue killings. But, I have argued, this defence fails: it is not generally true that being required to engage in a justified killing imposes a sufficiently serious cost on the rescuer's autonomy to render rescuing supererogatory. *Trolley* is, after all, a case in which we are permitted to impose lethal costs on an innocent person to prevent harm to the five. Even granting the moral significance of making Pedestrian harmfully treat herself as a means, the cost we propose for her cannot plausibly be thought to outstrip the cost we propose for Workman. The Permissibility Thesis must be rejected. In the

¹³ See Fabre (2007) for a defence of the view that one can be required to kill liable attackers in other defence.

following section, I consider the implications of the Requirement Thesis for the permissibility of Workman's defending himself in *Trolley*.

VI. DEFENCE AGAINST LESSER-EVIL HARMS

Here are three ways in which Workman might defend himself in *Trolley*. The first harms Pedestrian and prevents the saving of the five. The second harms Pedestrian without preventing the saving of the five. The third prevents the saving of the five without harming Pedestrian. As I will argue, the Requirement Thesis rightly prohibits defence in all of these cases.

The permissibility of Workman's defending himself seems least plausible in *Grenade*:

Grenade: Workman throws a grenade at the switch, killing Pedestrian as a side-effect. The trolley hits and kills the five.

McMahan argues that Workman's defence is impermissible because it will violate Pedestrian's right not to be killed. Ordinarily, someone who is morally responsible for an unjust threat of harm would render herself liable to defensive harm. But, McMahan argues, one cannot render oneself liable to harm when one threatens with objective justification, even if one was not required to act (2008; 2014: 105). Thus, even though Pedestrian will responsibly inflict an unjust harm on Workman, she does not thereby forfeit her right not to be harmed. The same is true in *Later Grenade*, when Workman's defence does not prevent the saving of the five:

Later Grenade: Pedestrian has diverted the trolley, saving the five. The trolley is now heading towards Workman. Workman can destroy the trolley by throwing a grenade that will also kill Pedestrian as a side-effect.

Here too, McMahan argues that Workman may not throw the grenade, since Pedestrian is not liable to be killed.

However, McMahan thinks that defence is permissible in *Jamming*:

Jamming: Workman uses a remote jamming device to prevent the switch from being flipped. Pedestrian is unable to divert the trolley and the five are killed.

McMahan gives two arguments for the permissibility of Workman's jamming the switch. First, Workman is not required to sacrifice himself for the five by turning the trolley towards himself. Given this, he may defend himself in a way that prevents the saving of the five (2009: 42; 2014: 110). Second, by jamming the switch, Workman does not inflict harm on the five. He only prevents their being saved. Savings we prevent count for less in our proportionality calculations compared to harms we cause (2014: 111). McMahan supports this argument with the following case, *Surgery*:

Surgery: Surgeon is trying to unjustly kill Victim. He will succeed unless Victim kills him. If Victim kills Surgeon, the five patients on whom Surgeon is scheduled to operate the following week will die.

It seems clearly permissible for Victim to kill Surgeon. If so, killing Surgeon cannot be disproportionate, despite the foreseen consequence of preventing the saving of the five. Similarly, McMahan claims, it is not disproportionate for Workman to save his own life in *Trolley* at the cost of preventing the saving of the five.

The Requirement Thesis, in contrast, gives us powerful reason to reject the permissibility of Workman's defending himself in all three cases. First, it gives us particularly strong reason to think that Pedestrian is not liable to be killed, since she is required (and not merely permitted) to divert. That Pedestrian has a duty to divert the trolley makes it especially hard to see how Workman might be permitted to prevent her from diverting. The facts that compel Pedestrian to divert are no less true for Workman. If they generate a duty for Pedestrian to act, they plausibly generate a duty for Workman not to interfere.

Moreover, it's particularly implausible that Pedestrian might become liable to harm by doing what morality compels her to do. It seems generally implausible that one might forfeit basic rights by doing one's duty, since this entails that morality might leave one with no morally permissible options for retaining one's rights. We might, perhaps, grant this possibility with respect to less important rights. But we should be sceptical that it obtains with respect to my right not to be physically harmed. Moreover, it seems straightforwardly incoherent that morality might pronounce Pedestrian liable to be harmed to prevent her from doing the very thing that morality requires her to do. The Requirement Thesis thus gives us a particularly compelling explanation of the wrongness of Workman's defending himself in *Grenade*.

Secondly, the Requirement Thesis holds that the five have claims to be rescued. This gives us grounds to prohibit Workman's defending himself in *Jamming* and supplements our grounds for such a prohibition in *Grenade*, since in both cases his defence prevents the saving of the five. It thus allows us to capture the worseness of Workman's defending himself in *Grenade* compared to *Later Grenade*, whereas explanations that focus solely on Pedestrian's liability do not distinguish these cases.

Prohibiting Workman's defence in *Jamming* conflicts with McMahan's view that Workman may defend himself. McMahan's conclusion regarding *Jamming* also looks to be in tension with the Requirement Thesis. If Workman is permitted to prevent the turning of the trolley, this challenges the idea that Pedestrian has a duty to turn it. However, I think we have good reason to reject both of McMahan's arguments.

McMahan's first argument—that Workman need not divert towards himself—conflates different ways in which one might be required make a

sacrifice for the sake of others, overlooking the moral significance of requiring a person to harmfully treat herself as a means. It is plausibly harder to justify requiring a person to harmfully treat herself as a means than requiring her to let herself be harmed as a side-effect. If so, it will not follow from the fact that Workman need not divert the trolley towards himself that he may prevent its being diverted.

McMahan's second argument overlooks a morally significant difference between *Surgery* and *Jamming*. It matters that if Victim is undefended in *Surgery*, she will suffer an *unjustified* killing. Surgeon is not killing Victim as some unfortunate but unavoidable side-effect of saving the five. Rather, his killing of Victim is wholly gratuitous. Requiring someone to let herself be subjected to gratuitous violence is very hard to justify because it is an especially egregious form of treating her as a means. Being harmfully treated as a means for the good of others is worse than being harmed as a foreseen side-effect of benefiting others.¹⁴

This is true even when the harm one allows a person to suffer is not gratuitous, as in *Bus*:

Bus: Child is about to fall under a bus. Pedestrian can grab her, saving her life. The bus will then hit five other children. If Pedestrian lets Child fall, the bus will kill her, but will then stop and not hit the other children (Tadros 2011: 122).

It is impermissible to lethally use Child as a means by intentionally letting her be hit by the bus in order to save the five. It is impermissible even though the killing of Child is instrumental in saving the five. It is plausibly even worse to use someone as a means by letting them suffer a wholly gratuitous harm, as we would do if we required Victim to let herself be murdered by Surgeon (that there is a reason for Victim not to defend herself does not, of course, entail that there is a reason for Surgeon to kill her. The harm she suffers is still gratuitous, even if failing to prevent it is not gratuitous.).

It is the unjustified, gratuitous killing of Victim, then, that is the harm that gets weighed against preventing the saving of the five in *Surgery*. Since unjustified, gratuitous harms are especially morally serious, we can see why it is proportionate for Victim (or a third party) to save Victim's life at the cost of preventing the saving of the five. But since being justifiably harmed as a side-effect is less morally bad than being unjustifiably and gratuitously harmed, it does not follow that Workman may prevent the trolley's being turned.

We can further undermine McMahan's argument by thinking more generally about when Victim would be permitted to defend herself against Surgeon. Very plausibly, Victim may prevent Surgeon's breaking her leg even if her

¹⁴ There is an extensive literature on this topic that I cannot explore here, but the worseness of harmfully using compared to harming as a foreseen side-effect is widely accepted amongst non-consequentialists. See, e.g. Tadros (2011, 2015: 117–22).

defence would render Surgeon unfit to operate (imagine that she breaks Surgeon's wrist). But I doubt McMahan would want to allow that, therefore, Workman may prevent the diverting of the trolley if the trolley would break his leg. If so, the permissibility of defence against unjustified, gratuitous harms does not extend to defence against equivalent justified harms. McMahan has not shown that Workman may defend himself in *Jamming*.

On the contrary, the reasons we have to accept the Requirement Thesis give us good grounds to prohibit Workman's defence in *Jamming*. According to *Preventing Harm*, Pedestrian has a duty to divert the trolley. This duty is owed to the five, who have correlating claims to be saved, since they can be saved without unreasonable cost to Pedestrian and without violating anyone's rights—that is, without anyone's being unjustifiably harmed. Preventing the trolley's being turned therefore violates the five's claims to be saved. In contrast, the five in *Surgery* have no claim that Victim let herself be unjustifiably and gratuitously killed if so that they are saved. Victim's defence does not, therefore, violate their claims.

One might argue that defence in *Later Grenade* is in fact permissible, on the following grounds. Once Pedestrian has diverted the trolley, and the rights of the five are no longer at stake, this case is best understood in terms of who ought to bear the cost of the rescue. And, there's no reason why these costs should fall on Workman rather than Pedestrian (Quong 2016: 827). Perhaps Workman is permitted to make Pedestrian bear the cost of the rescue if, for example, he does so only after a fair procedure such as tossing a coin.

On this view, Pedestrian poses a sort of innocent threat to Workman. But, as I have argued elsewhere, the permissibility of using defensive force against innocent threats is partly explained by the fact that even innocent people have duties not to harm others and can be made to bear significant costs to comply with those duties (Frowe 2014: 67–70). But, by hypothesis, Pedestrian has no such duty not to harm Workman. Killing Pedestrian does not, therefore, permissibly force her to comply with a duty. Rather, killing Pedestrian is akin to killing a bystander who is under no duty to bear significant costs to save Workman's life. I do not think that one may toss a coin to decide whether to kill a bystander to save one's own life. If that is correct, this argument does not show the permissibility of Workman's killing Pedestrian.

Note that the most influential alternative accounts of a permission to kill innocent threats—namely, those grounded in agent-relative reasons (Davis 1984; Fabre 2012; Quong 2009, 2016)—also fail to undermine the Requirement Thesis. At best, these arguments show that Workman has a special permission to override a general prohibition on killing Pedestrian.¹⁵ It does not show that Pedestrian lacks a duty to divert the trolley.

¹⁵ Quong (2016) defends the permissibility of Workman's killing Pedestrian on these grounds.

VII. CONCLUSION

I have argued that we should explicitly recognise cases in which we have lesser-evil justifications for harming for the sake of others as rescue cases. Ordinarily, one may refrain from justified rescues only if rescuing requires one to bear a supererogatory cost. And yet it is widely believed that acting on lesser-evil justifications is morally optional, unless there is some gross disparity of harms. The most common defence of this optionality is the cost we impose on a rescuer's autonomy by requiring her to inflict harms on innocent people. But, as I have argued, it is not plausible to suppose that the cost to a rescuer's autonomy ordinarily renders lesser-evil rescues supererogatory. I have therefore defended the Requirement Thesis, which holds that unless rescuing imposes a very serious psychological cost on the rescuer, of the sort that will seriously impede her capacity for flourishing, she is required to enact lesser-evil rescues. If saving five lives justifies killing an innocent person, it also justifies requiring a rescuer to do something she does not want to do. The final part of the paper argued that the Requirement Thesis gives us a plausible account of why Workman may not defensively harm Pedestrian, nor prevent the saving of the five, that improves upon accounts that focus solely on Pedestrian's liability.¹⁶

REFERENCES

- Alexander, L. (2005) 'Lesser-Evils: A Closer Look at the Paradigmatic Justification', *Law and Philosophy*, 24: 611–43.
- Bazargan, S. (2016) 'Killing Minimally Responsible Threats', *Ethics*, 125: 114–36.
- Davis, N.A. (1984) 'Abortion and Self-Defence', *Philosophy and Public Affairs*, 13: 175–207.
- Fabre, C. (2007) 'Mandatory Rescue Killings', *Journal of Political Philosophy*, 15: 363–84.
- (2012) *Cosmopolitan War*. Oxford: OUP.
- Feinberg, J. (1972) 'Voluntary Euthanasia and the Inalienable Right to Life', *Philosophy and Public Affairs*, 7: 93–123.
- FitzPatrick, W.J. (2009) 'Thomson's Turnabout on the Trolley', *Analysis*, 69: 636–43.

¹⁶ Drafts of this paper were presented at the Moral Philosophy Seminar, Oxford University; Law and Philosophy Seminar, Jindal Global University; Westmont College, Santa Barbara; Workshop on Self-Defence, University of Southern California; the Philosophy Department Colloquium, Stockholm University; the Society for Applied Philosophy Annual Conference, Edinburgh; Law School Seminar of the Ethics and Law of Armed Conflict, Harvard University; Joint Philosophy and Political Science Colloquium, UC Boulder; Centre for Ethics, University of Toronto; Annual Graduate Philosophy Conference, Western Michigan University; Philosophical Society of Southern Africa, Fort Hare; Department Colloquium, Antwerp University; Department Colloquium, Yale Law School. I am very grateful to participants at all of these events. Special thanks to Adam Hosein, Lars Christie, Jonathan Quong, Massimo Renzo, Victor Tadros and Alec Walen for extremely helpful comments and discussion. I am especially grateful to the two anonymous referees for this journal, whose wonderfully constructive criticisms significantly improved the paper. This research was funded by the Knut and Alice Wallenberg Foundation, and further supported by a fellowship at the Jack and Mae Nathanson Centre for Transnational Human Rights, Crime and Security at York University, Canada.

- Foot, P. (1967) 'The Problem of Abortion and the Doctrine of Double Effect', *Oxford Review*, 5: 1–7.
- Frowe, H. (2008) 'Equating Innocent Threats and Bystanders', *Journal of Applied Philosophy*, 24: 277–90.
- (2014) *Defensive Killing*, Oxford: Oxford University Press.
- (2015) 'Can Reductive Individualists Allow Defence Against Political Aggression?', in P. Vallentyne, D. Sobel and S. Wall (eds) *Oxford Studies in Political Philosophy*, vol. 1, 173–93, New York: OUP.
- Hallborg, R., Jr. (1997) 'Comparing Harms: The Lesser-Evil Defense and the Trolley Problem', *Legal Theory*, 3: 291–316.
- Kagan, S. (1989) *The Limits of Morality*, Oxford: Clarendon Press.
- Kamm, F.M. (2007) *Intricate Ethics: Rights, Responsibilities and Permissible Harm*, Oxford: OUP.
- Lazar, S. (2014) 'National Defence, Self-Defence and the Problem of Political Aggression', in C. Fabre and S. Lazar (eds) *The Morality of Defensive War*, 11–39, Oxford: OUP.
- (2015) *Sparing Civilians*, Oxford: OUP.
- MacAskill, W. (2015) *Doing Good Better: Effective Altruism and a Radical New Way to Make a Difference*, London: Faber & Fabre.
- McMahan, J. (2008) 'Debate: Justification and Liability in War', *Journal of Political Philosophy*, 16: 227–44.
- (2009) *Killing in War*, Oxford: OUP.
- (2014) 'Self-Defence Against Justified Threateners', in H. Frowe and G. Lang (eds) *How We Fight: Ethics in War*, 104–37, Oxford: OUP.
- Moore, M. (2014) 'Collective Self-Determination, Institutions of Justice, and Wars of National Defence', in C. Fabre and S. Lazar (eds) *The Morality of Defensive War*, 185–202, Oxford: OUP.
- Oberman, K. (2015) 'The Myth of the Optional War: Why States Are Required to Wage the Wars They Are Permitted to Wage', *Philosophy and Public Affairs*, 43: 255–86.
- Parfit, D. (2011) *On What Matters*, Oxford: OUP.
- Pummer, T. (2017) 'Whether and Where to Give', *Philosophy and Public Affairs*, 44: 77–88.
- Quong, J. (2009) 'Killing in Self-Defence', *Ethics*, 119: 507–37.
- (2015) 'Rights Against Harm', *Proceedings of the Aristotelian Society, Supplementary Volume*, 89: 249–66.
- (2016) 'Agent-Relative Prerogatives to Do Harm', *Criminal Law and Philosophy*, 10: 815–29.
- Rodin, D. (2011) 'Justifying Harm', *Ethics*, 112: 74–110.
- (2017) 'The Lesser Evil Obligation', in S. Bazargan and S.C. Rickless (eds) *Ethics in War*, 28–45, New York: OUP.
- Tadros, V. (2011) *The Ends of Harm: The Moral Foundations of Criminal Law*, Oxford: OUP.
- Tadros, V. (2015) 'Dimensions of Intentions: Ways of Killing in War', in S. Lazar and H. Frowe (eds) *Oxford Handbook of Ethics of War*, New York: Oxford University Press, DOI 10.1093/oxfordhb/9780199943418.013.4.
- Thomson, J.J. (1985) 'The Trolley Problem', *Yale Law Journal*, 94: 1395–415.
- (1986) 'Some Ruminations on Rights', in W. Parent (ed.) *Rights, Restitution and Risk: Essays in Moral Theory*, 49–65, Cambridge, MA: Harvard University Press.
- (2008) 'Turning the Trolley', *Philosophy and Public Affairs*, 36: 359–74.
- and D. and Wasserman, (2012) 'Agents, Impartiality and the Priority of Claims Over Duties: Diagnosing Why Thomson Still Gets the Trolley Problem Wrong by Appeal to the "Mechanics of Claims"', *Journal of Moral Philosophy*, 9: 545–71.
- Walzer, M. (1977) *Just and Unjust Wars: A Moral Argument with Historical Illustrations*, New York: Basic Books.

Stockholm University, Sweden