

II—HELEN FROWE

CLAIM RIGHTS, DUTIES, AND LESSER-EVIL JUSTIFICATIONS

This paper explores the relationship between a person's claim right not to be harmed and the duties this claim confers on others. I argue that we should reject Jonathan Quong's evidence-based account of this relationship, which holds that an agent *A*'s possession of a claim against *B* is partly determined by whether it would be reasonable for *A* to demand *B*'s compliance with a correlative duty. When *B*'s evidence is that demanding compliance would not be reasonable, *A* cannot have a claim against *B*. I suggest that some of the putatively problematic cases that Quong identifies can be resolved by plausibly narrowing the scope of the right not to be harmed. I also argue that Quong's view leads to implausible conclusions, and his account of what happens to *A*'s claim in the face of lesser-evil justifications is inconsistent with his broader view. I then defend the view that agents are required, and not merely permitted, to act on lesser-evil justifications. I further argue that *A* may not defend herself against the infliction of harms that are justified on lesser-evil grounds. However, she may defend herself in cases where *B* is only evidentially, and not objectively, justified in harming her.

I

Quong's Account. Deontological ethical theories usually hold that people have rights of various sorts. The right not to be harmed is usually described as a claim right (henceforth, following Quong 2015, a claim). Claims are distinctive because they confer both duties of non-interference and duties of assistance on other people. If *A* has a claim that *B* not harm her, this entails both that *B* has a duty to refrain from harming *A* and that third parties ought to assist *A* if *B* tries to harm her, assuming that doing so is not too costly for them.

How do these thoughts about claims and duties play out in a case such as *Day's End*?

Day's End: *B* always comes home at 9:00 p.m., and the first thing he does is to flip the light switch in his hallway. He did so this evening. *B*'s flipping the switch caused a circuit to close. By virtue of an extraordinary series of coincidences, unpredictable in advance by anybody, the circuit's closing caused a release of electricity (a small lightning flash) in *A*'s house next door. Unluckily, *A* was in its path and was therefore badly burned. (Thomson 1990, p. 229)

According to Judith Thomson (1990), we each possess certain fundamental claims, such as the claim not to be harmed, in virtue of being human. Thomson defends a fact-relative view of when agents' claims are violated, denying that one needs fault or even agency to violate a right. All that is required for *B* to violate *A*'s right is that *A* has a right not to be harmed, and *B* harms *A*. *A* has done nothing to dispossess herself of her claim in *Day's End* (she has not, for example, waived or forfeited her claim), and so in harming her, *B* violates her claim.

Quong defends a largely evidence-relative view of when agents have claims. According to Quong, *A* has a claim against *B* only if it is reasonable for *A* to demand *B*'s compliance with the correlative duties.¹ Quong argues that reasonableness is a moralized notion that extends beyond the mere cost of compliance. Reasonableness is sensitive to 'the interests and status of all affected parties', and includes 'questions about the costs of refraining and the opportunities to avoid refraining' (Quong 2015, pp. 255, 257). And the considerations that determine the reasonableness of *A*'s demand must be assessed from the perspective of the potential duty-bearer, *B*, with reference 'to abilities and information that can reasonably be expected to be available to [*B*] in a given situation' (Quong 2015, p. 257).

Quong argues that, from *B*'s perspective, *A* cannot reasonably demand that *B* refrain from harming her in *Day's End*. Given that *B* can never be certain that flipping a light switch will not harm anyone, a duty to refrain from harming *A* would amount to a requirement that *B* never turn on a light switch (or engage in any other potentially harmful behaviour). Compliance with such a duty is too burdensome to be reasonable, and *B* therefore has no duty to refrain from flipping the switch (Quong 2015, p. 258). If claims consist in their ability to impose duties on others, *A* can have no claim not to be harmed in *Day's End*. It follows, of course, that *B* does not vio-

¹ Compare with Joseph Raz's claim that *X* has a right if, and only if, 'an aspect of *X*'s well-being (his interest) is a sufficient reason for holding some other person(s) to be under a duty' (Raz 1984, p. 195).

late any claim of *A*'s in harming her.

Quong supports his analysis by pointing out that our claim to be aided is explicitly dependent upon cost. When saving a person is too costly, we do not violate or infringe her claim if we fail to save her. Rather, he argues, she simply lacks a claim to be aided altogether. Quong's thesis is that claims in general are like this. *B* has a duty not to ϕ with respect to *A* only if it is reasonable for *A* to demand that *B* not ϕ . Part of what determines the reasonableness of such a demand is the cost of refraining from ϕ -ing. When it is too costly for *B* to refrain, the demand is not reasonable and *B* had no duty not to ϕ .

II

Reasonableness, Cost, and the Right not to be Harmed. As noted above, Quong draws a distinction between reasonableness and mere cost, arguing that the mere cost to *B* of refraining from acting cannot determine whether *B* is under a duty to so refrain. He illustrates this by comparing two cases, which I'll call *Transplant* and *Naughty Transplant*.

Transplant: *A* will die unless she receives a kidney transplant. *B*, who also needs a transplant to survive, arrives at the hospital before *A*, and is therefore placed above *A* on the transplant waiting list. *B* receives a kidney and survives. *A* does not, and dies.

Naughty Transplant: *B* will die unless he receives a kidney transplant. The only available kidney is *A*'s, so *B* lethally takes *A*'s kidney.

In both cases, the cost to *B* of refraining from acting is the same—he will die if he does not arrive at the hospital first, and he will die if he does not take *A*'s kidney. And the cost to *A* is the same as well—*A* will die if *B* arrives first, and she will die if *B* takes her kidney. And yet, it seems that *A* has a claim that *B* refrain from acting in *Naughty Transplant*, and not in *Transplant*. If cost were doing all the work, we would be unable to explain this difference. Quong therefore concludes that any account of when agents have claims must include something like his reasonableness criterion, which is based on: '(i) the value of the activity *B* engages in by ϕ -ing, (ii) how costly it would be for *B* to refrain from ϕ -ing, (iii) whether *B* had sufficient

opportunity to avoid finding himself in situation where he faces the choice whether or not to ϕ , and (iv) whether *A* had sufficient opportunity to avoid finding [herself] in situation where [s]he might be harmed by *B*'s ϕ -ing' (Quong 2015, p. 255).

It's not clear how these considerations distinguish *Transplant* from *Naughty Transplant*. The value of the activity seems the same in each case—acting saves *B*'s life, but costs *A*'s life. We've already agreed that the cost to *B* of refraining is the same in each case. Let's assume that in *Naughty Transplant*, *B* had no control over developing the illness that means he needs a new kidney, and that *A* had no control over being in possession of a functioning kidney. This means that there are no opportunities to avoid their positions vis-à-vis each other, so neither can have unreasonably failed to take such an opportunity. Both are innocent, so invoking their moral status won't help, if by that we mean that, for example, one is at fault or liable to some harm. Similarly, given their equal innocence, there's no obvious reason to prefer the interests of *A* to the interests of *B*, once we ignore *A*'s claim that people not take her kidney. And, of course, we must ignore *A*'s claim that people not take her kidney, because whether she has such a claim is precisely what the reasonableness test is meant to establish. So Quong's view fails to give us a way to distinguish these cases.

How else might we explain the difference between *Transplant* and *Naughty Transplant*? I think the apparent difficulty here arises only because Quong is construing the claim not to be harmed in an implausibly broad way. He writes as if all instances of harming *A* prima facie fall under the remit of *A*'s claim not to be harmed. A narrower, more plausible interpretation of this claim holds that it is a claim against people harmfully interfering with our bodies or with things to which we are entitled, where such interference results in physical harm to us.² If we restrict the remit of the claim not to be

² I'm including interference with tools here—you count as interfering with my body if you poke me with a stick. You also count as interfering with my body if, for example, you switch my painkillers for poison. But turning up early at the hospital does not harm me by interfering with my body. It only stops me from being able to get what I need. When I have a claim over the thing you tamper with, this can count as violating my right not to be harmed. But that's because I already have an entitlement to the pills, whereas we're assuming (with Quong) that entitlement to the kidney lies with whoever turns up first. (Of course, it's possible that entitlement to the kidney should be determined on some other basis—perhaps if *A* will get the greatest benefit, she should get the kidney. But then, by using the kidney in *Transplant*, *B* would be taking something to which *A* is entitled independently of who turns up first. On my account, that would bring *B*'s using of the kidney under the remit of *A*'s right not to be harmed.)

harmed in this way, we can explain the difference between the two *Transplant* cases quite easily. When *B* arrives at the hospital ahead of *A*, he does not interfere with *A*'s body, nor with anything to which *A* is entitled (since, we're assuming, entitlement to the kidney depends upon arriving first). It doesn't matter, therefore, that *B* is causing harm to *A*, or what it would cost him to avoid causing this harm, because not all ways of harming *A* fall under the remit of *A*'s claim not to be harmed. *Naughty Transplant*, in contrast, does involve harmfully interfering with *A*'s body, and so it potentially falls under the scope of *A*'s claim not to be harmed.

A fact-relative account of claims, such as Thomson's, will hold that since *A* has done nothing to lose her claim in *Naughty Transplant*, *B* violates her claim in taking her kidney. An evidence-relative account, such as Quong's, could invoke the distinction between doing and allowing to show that *B* violates *A*'s claim. The cost to *B* of refraining is not greater than the cost to *A* if *B* acts. If harms we cause count for more, morally speaking, than the costs we allow, this gives us an asymmetry between *B*'s bearing the cost and *A*'s bearing the cost. This could make it reasonable for *A* to demand that *B* comply with a duty not to take her kidney.

III

Evidence and Claims. This narrower account of the claim not to be harmed resolves some of the cases that Quong discusses in a way that is compatible with both evidence-relative and fact-relative accounts of when agents have claims. But there are cases in which Quong's evidence-relative account of which claims we possess is deeply counter-intuitive. Consider *Duped Soldiers*:

Duped Soldiers: A group of young soldiers are fooled by a totalitarian regime into believing that their regime is good and just, and is being attacked by its evil neighbours, the Gloops. The regime's misinformation campaign is subtle and absolutely convincing: the soldiers are justified in believing what they are told. The soldiers are given orders to attack a Gloop village that, they are told, is a terrorist camp where Gloops are plotting a major attack. The inhabitants of the village are in fact innocent civilians who will be killed unless a neutral peacekeeping force intervenes to kill the soldiers. (Quong 2015, p. 261)

Quong says that it would be ‘unacceptable’ to think that if the soldiers have sufficient evidence that the villagers are terrorists, the villagers then have no claims against being harmed by the soldiers (Quong 2015, p. 261).

He tries to show that this is not an implication of his argument by distinguishing two questions:

- (i) Under what conditions does a person have a claim not to be harmed by a particular type of act performed by another person?
- (ii) Has some particular person, *A*, waived, transferred or forfeited this claim not to be harmed by another person, *B*?

(Quong 2015, p. 261)

He argues that (i)—whether one has a claim right to begin with—depends on what it is reasonable for *A* to demand of *B*. And he thinks that there are certain *types* of actions from which one can reasonably demand that others refrain, ‘taking various factors into account, including the obviously very high chance that [acting] will cause serious harm or death’ (Quong 2015, p. 262). One type of action from which one can reasonably demand that other people refrain is shelling one’s home. Thus, we can establish initial claims on the part of the villagers. And the second question—whether the villagers have done anything to dispossess themselves of those claims—is not answered with reference to the soldiers’ evidence, but is rather determined objectively by looking at what the villagers have actually done. Since they have not transferred, waived or forfeited their claims, they have a claim not to be killed by the soldiers.

It’s not clear how Quong can endorse an objective account of the dispossession of claims. If the soldiers’ evidence is that the villagers have forfeited or waived their claims, why would it be reasonable to demand the soldiers’ compliance with the correlative duties? But even setting that concern aside, I don’t think this response is convincing, for two reasons.

First, identifying the types of acts from which one can reasonably demand that others refrain seems quite open. The soldiers are shelling a house, but they are also engaging in what their evidence suggests is legitimate self- and other-defence. So, for this idea to be helpful, we need a further argument for why ‘shelling one’s house’ is the relevant description for determining what type of act the soldiers are engaged in.

We might think that Quong is trying to draw a distinction be-

tween cases in which I know I pose a threat, but am mistaken about whether that threat is just, and cases in which I don't know that I pose a threat at all. He says, for example, that

When the type of act that *B* performs is one that *A* cannot reasonably demand he refrain from performing—say, because it does not pose any undue risk of harm and refraining would be very costly or damaging to *B*—then *A* has no claim that *B* refrain. This is true even if *A* would have such a claim if *B*'s epistemic situation were different. (Quong 2015, p. 259)

He also argues that the soldiers ought to 'understand that the permissibility of their proposed course of action hinges on their assumption that the [villagers] have forfeited their rights' (Quong 2015, p. 262). But the point of the case is that the soldiers do indeed have grounds for such an assumption: they are *Q* in their beliefs that the villagers are dangerous terrorists, and that not shelling is unreasonable. This entails that the soldiers' evidence defeats the usual presumption against shelling people's houses. Their evidence is that their actions do not impose an *undue* risk of harm because shelling is justified *despite* the fact it is intentionally harmful.

Second, it's hard to see how the type of action, rather than the token, could be relevant on an evidence-relative account.³ Quong's argument asserts a claim prior to establishing the reasonableness of demanding compliance in any token case. But if one is proposing an evidence-relative account to determine the existence of claims, one must base the existence of a claim on the evidence in any token case. At best, the general reasonableness of a type of action might give us a standard by which we can judge a person's evidence for acting in a token case—the standard of evidence required before acting when one knows that one will cause harm could be more stringent than in cases where one has no reason to suppose one will cause harm. But if there is sufficient evidence that not shelling the house is unreasonable in a token case, then the fact that shelling is usually wrong can be no bar at all to shelling in this token case on an evidence-relative account.

We can see this clearly by thinking about *Day's End*. The implication of Quong's argument is that *A* has no claim that people refrain from the type of action that is turning on their light switches, be-

³ Thanks to Seth Lazar for suggesting that I frame the problem in this way.

cause this type of action does not pose an undue risk. Hence, *A* seemingly lacks a claim in the way Quong suggests. But imagine that *B*'s evidence in a token case is that flipping the light switch will harm *A*. *B* surely cannot point to the lack of a claim regarding the type of act that is flipping a light switch to show that *A* has no claim when the evidence suggests that a token flipping is harmful. Rather, given the evidence, *A* has a claim against the token flipping irrespective of how reasonable it is to demand that *B* refrain from flipping light switches in general.

If Quong grants *A* a claim in this case, he is committed to the view that it is the evidence concerning any token act that matters for establishing claims, not a kind of general reasonableness about demanding compliance with the type of act. And this should be similarly true in *Duped Soldiers*. Even if we were to grant the villagers a general claim not to be subjected to the type of act that is shelling, they must lack this claim in any token case where the soldiers' evidence is that not harming the villagers will be unreasonable. That conclusion is, as Quong says, unacceptable.

IV

Lesser-Evil Justifications. Quong takes a different approach to claims in cases where *B* has a lesser-evil justification for harming *A*. Lesser-evil justifications obtain when one will prevent substantially more harm than one causes, such that the disparity between the harm and the good overrides the deontological presumption against causing harm. They are also sensitive to modes of agency, such as whether one will be causing harm as a side effect rather than as a means.

A lesser-evil justification for harming is generally thought to obtain in *Trolley*:

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

Most people think that the disparity between the harm inflicted on *A* and the good secured by saving the five is sufficient to give *B* a lesser-evil justification for diverting the trolley. Most people also

think that in diverting the trolley, *B* will permissibly infringe, but not violate, *A*'s claim not to be harmed. This of course entails that *A* has a claim not to be harmed. This verdict is prima facie incompatible with Quong's view. *B*'s evidence is that not harming *A* is too costly to be reasonable, and thus *A* should lack a claim that *B* refrain from harming her.

Quong tries to resist this result by arguing that *A* can have a claim against *B* when demanding the claim's fulfilment *would be* reasonable were it not for the existence of a lesser-evil justification. *B* would have a duty not to kill *A* if the five were not on the other track. Quong's thesis thus holds that lesser-evil justifications outweigh, but do not negate, claims (Quong 2015, p. 9 n.12 <INT REF>) This means that there are, after all, cases in which *A* may not reasonably demand that *B* not harm her, and yet *A* retains a claim not to be harmed.

This caveat causes at least two problems for Quong's account. The first is that it seems to undermine the distinctive status of a claim by denying that *A*'s claim confers a duty on *B*. Given the presence of the five, it is not reasonable to demand that *B* refrain from harming *A*. If there is no duty to refrain when a demand to refrain is unreasonable, and claims require duties, *A* can have no claim in *Trolley*.

Second, if Quong takes this approach in *Trolley*, where *A* has a claim even though demanding compliance is unreasonable, it seems it should also apply in *Day's End*. *A* has a claim not to be harmed, but given that it's unreasonable to expect *B* to know that his action will harm *A*, it's unreasonable to demand *B*'s compliance.

Quong might revise his argument to suggest that *B* has a *pro tanto* duty not to harm *A* in *Trolley* that correlates with *A*'s claim, but that this duty is overridden by the duty to save the five.⁴ Thus the case is unlike *Day's End*, because in *Day's End* *A* simply lacks a claim altogether. But I don't think Quong's argument from reasonableness supports this way of distinguishing the cases. His argument is that *B* cannot be under a duty to refrain from harming *A* if demanding *B*'s compliance with that duty is unreasonable. Here, demanding *B*'s compliance is unreasonable, and so there is no scope for grounding even a *pro tanto* duty. Indeed, I'm not sure that Quong's view can even make sense of merely *pro tanto* duties (as opposed to all-things-

⁴ Thanks to Massimo Renzo for this suggestion.

considered duties), since one cannot reasonably demand compliance with merely *pro tanto* duties. Furthermore, most people deny that there is a duty to save the five. The popular view is that *B* is merely permitted to save the five. It's hard to see how a mere permission could outweigh a genuine duty. I return to this point in §V below.

Once one grants that claims can persist even when demanding compliance is unreasonable, it's not clear why *A*'s claim persists when *B* is both evidentially and objectively justified in acting (as in *Trolley*), but is negated when *B* is merely evidentially justified in acting (as in *Day's End*). Yet Quong's view generates this peculiar result. If, on the one hand, *B*'s evidence is what matters, *A* should lack a claim in both *Trolley* and *Day's End*. *B*'s evidence in both cases is that refraining from acting is too costly to be reasonably demanded by *A*. If, on the other hand, the fact that diverting the trolley is objectively justified is allowed a role in determining whether *A* has a claim, it should surely push in the opposite direction to that endorsed by Quong. *A* should lack a right when there's objective justification for acting, and retain a right when there's only an evidenced-based justification for acting.

V

Turning the Trolley. So, Quong's account of *Trolley* is unpersuasive. But the mechanics of lesser-evil justifications are somewhat mysterious, and I don't think deontologists have yet given a satisfactory account of them.

In this section, I argue that we should reject the commonly held view that one can be permitted but not required to act on a lesser-evil justification.⁵ Only when things are bad enough that one has a *duty* to kill an innocent person to save others could it be *permissible* to kill an innocent person. Call this the Requirement Thesis. It contrasts with the Permissibility Thesis (that diverting is merely permissible) and with the Impermissibility Thesis (that diverting is impermissible). Judith Thomson has lately defended the Impermissibility Thesis, but it's not a view that is widely endorsed, so I won't address it here.

⁵ Thanks to Massimo Renzo and Jeff McMahan for helpful discussion of this idea.

The Requirement Thesis. It's hard to identify the precise number of lives at which point saving is required even if it involves killing one. I think the moral significance of the distinction between doing and allowing is very weighty, and that perhaps five lives are not enough to overcome the presumption against doing harm. However, as we increase the number of lives to be saved such that diverting the trolley becomes clearly permissible, I think we thereby make diverting plausibly required. I should think that ten lives is enough to require diverting, but there might be an inevitable vagueness about whether eight or nine is also sufficient.⁶

I suspect that it is in part this vagueness that has attracted people to the idea that the moral status of diverting moves from impermissible to permissible to required as we increase the number of lives to be saved from one to ten. Mere permissibility might seem like a kind of middle ground that is an appropriate response to uncertainty about when the lesser-evil justification kicks in.

But that's a mistake that reveals the central motivating thought of the Requirement Thesis. If we reject the Requirement Thesis in favour of the Permissibility Thesis, we seem to be committed to the idea that the one who is in a position to rescue gets to decide whether or not to rescue based on his or her preferences. That strikes me as an deeply unattractive position: it's morally inappropriate to think that whether an innocent, non-threatening person like *A* gets to live depends on whether *B* wants to kill her. We shouldn't think of mere permissibility as a happy middle ground in the face of such uncertainty. The appropriate response to uncertainty about the justifiability of killing *A* is refraining from killing her.⁷ It is not to leave killing her to *B*'s discretion.

Alec Walen and David Wasserman deny this, arguing that killing *A* should be left to *B*'s discretion, because this 'serves to protect an agent from having to treat herself as a mere tool for the promotion of the general welfare' (Walen and Wasserman 2012, p. 554). But this argument will succeed only if requiring agents to act on lesser-evil justifications to inflict lethal harm fails to leave agents an adequate sphere in which they need not treat themselves as tools. After all, it is not as if requiring agents to act on lesser-evil justifications

⁶ However, I'll present the cases as involving five lives since this is standard in the literature.

⁷ I think we should distinguish between cases where we know that if certain empirical facts obtain we are morally justified in harming and cases in which we know the empirical facts but are not sure about whether, given these facts, a moral justification obtains.

would require agents to treat themselves as tools *whenever* doing so would maximize the good. It would merely require them to do so sometimes. And this is surely plausible. If diverting the trolley will only break both *A*'s legs rather than kill her, *B* cannot invoke a claim not to treat himself as a tool to justify not diverting the trolley.

Given this, the question is *how much* we may restrict the sphere in which an agent need not treat herself as a tool, not *whether* we may so restrict it. And if *B* must divert when doing so will merely break *A*'s legs, there's reason to think that the issue is simply one of degree. Perhaps many lives must be at stake for *B* to be required to divert when he will kill, but it's not obvious why there would be a threshold at which point increasing the number of lives *B* could save has no bearing on whether he is required to act. Moreover, even this proposed restricting of the agent's sphere of discretion will be mitigated if, as I have suggested, we ought to revise our judgement about how many lives must be at stake to justify diverting the trolley. If a lesser-evil justification obtains only when, say, ten lives can be saved, requiring agents to treat themselves as mere tools in the face of lesser-evil justifications is unlikely to constitute a significant further erosion of the agent's protection against having to treat herself in this way.

Exceptions? It might seem that the following sort of case is an exception to the Requirement Thesis.

Grenade: Attacker is going to kill Victim. Victim can defend himself only by throwing a grenade at Attacker. However, the blast from the grenade will inflict mild burns on an innocent third party, Neighbour.

It seems plausible that Victim has a lesser-evil justification for defending himself foreseeing the harm he will cause to Neighbour. But it also seems plausible that Victim does not *have* to defend himself. Rather, he is permitted to choose to let himself be killed rather than harm Neighbour. Thus, the Requirement Thesis looks false in cases where one has a lesser-evil justification for saving oneself.

I think we can defend the Requirement Thesis against this objection. When Victim is deciding whether to defend himself, he can decide to discount harm to himself.⁸ But once we discount the good of

⁸ See Parry (forthcoming) for a discussion of how consenting to being harmed, and refusing consent to being rescued, affects proportionality.

saving Victim's life, saving him will not outweigh the harm done to Neighbour. The lesser-evil justification will therefore no longer obtain.

We can say something similar in the following case:

Legs: *B* can divert the trolley, killing the one and saving the five, only at the cost of his own legs.

Let's say that *B* isn't required to inflict the loss of his legs on himself, and so he isn't required to divert the trolley. (This view assumes that it is especially demanding to require an agent to inflict harm on herself, and that this explains why *B* may make *A* bear a lethal cost even if he need not make himself bear a less serious cost. But proponents of the Permissibility Thesis will presumably endorse the suggestion that *B* need not divert in this case, since they think he need not divert when doing so will not harm him at all.) Nonetheless, it seems that *B* would be permitted to divert: he could make a supererogatory sacrifice of his leg for the good of the five. And he's permitted to do so only because there's a lesser-evil justification for killing the one. Thus *Legs* looks like a case in which one is permitted but not required to act on a lesser-evil justification.

Here too I think that if *B* does not consent to inflicting the loss of his leg upon himself, there is no lesser-evil justification for diverting the trolley. If we're granting that inflicting the loss of his legs on himself is too demanding to be required, letting the five die will be the lesser evil. But if *B* consents, we can discount the harms to him in the way described above, and then (but only then) a lesser-evil justification obtains and he is required to divert.

Notice, then, that this case is not guilty of the moral inappropriateness I identified above. In supererogatory cases, whether the lesser-evil justification obtains is determined by *B*'s preferences (or his consent). The moral inappropriateness objection holds only against the view that *B* has discretion concerning whether to act on a lesser-evil justification that obtains *independently* of his preferences.

It might seem that this reliance on *B*'s consent results in a funny kind of duty, since the existence of the duty depends on *B*'s doing what is supererogatory. But duties can indeed be conditional in this way. As Shelley Kagan has argued, rushing into a burning building at the risk of one's own life is an act of supererogation. But if one does rush in, one is duty-bound to save the child trapped inside instead of the caged bird (Kagan 1989, p. 4).

VI

Self-Defence. One question that arises if we think that *B* has a duty to turn the trolley is what, if anything, *A* may do to try to stop him. Here are three ways in which *A* might defend herself against being killed by the trolley:

- (i) Harm *B* to prevent his diverting the trolley.
- (ii) Harm *B* after he has diverted the trolley when doing so will not prevent the saving of the five. (Imagine that the trolley has been diverted, so the five are saved, but it is now heading towards *A*. *A* can destroy the trolley by throwing a grenade that will also kill *B* as a side-effect.)
- (iii) Prevent *B*'s turning the trolley in a way that doesn't harm *B*. (Imagine *A* has a jamming device that she can use to remotely prevent the switch from being pulled.)

I think that all three courses of action are impermissible. *B* has no duty not to divert the trolley—this is true even if one thinks diverting merely permissible. And when there is no lesser-evil justification, defensive force is permissible only if it's a way of getting people to comply with their duties.⁹ Hence, using defensive force against *B* as in (i) and (ii) is impermissible.

We might reply that once *B* has diverted the trolley, as in (ii), the 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 *A* rather than *B*.¹⁰ Invoking the distinction between killing and letting die cannot help, since each will kill the other. So perhaps *A* is permitted to make *B* bear the cost of the rescue if, for example, she does so only after a fair procedure such as tossing a coin.

I think that *B* poses an innocent threat to *A*. I've argued elsewhere that defensively killing innocent threats is sometimes permissible. But, as described above, that permissibility is partly explained by the fact that even innocent people have duties not to harm others and can be made to bear serious costs to comply with those duties. But, by hypothesis, *B* had no such duty in *Trolley*. Killing *B* does not, therefore, permissibly force him to comply with a duty. Rather,

⁹ See Tadros (2011) for a sustained defence of this view.

¹⁰ Thanks to Massimo Renzo for this suggestion.

killing *B* is akin to killing a bystander, who is under no duty to bear very serious costs to save *A*'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. And so, I still think that (ii) must be impermissible.

I also think that using non-forceful methods to prevent *B*'s diverting the trolley, as in (iii), is impermissible. It is implausible that the facts could generate a duty for *B* to divert, whilst also generating a permission for *A* to prevent *B* from diverting. But Jeff McMahan rejects the claim that *A* may not prevent the diverting of the trolley.¹¹ McMahan claims that if *A* is not required to sacrifice herself, she may defend herself in a way that prevents the saving of the five.

But I think there's a morally significant difference between requiring that *A* sacrifice herself—that is, turn the trolley towards herself—and refusing her a permission to prevent the trolley's being turned. This difference is grounded in the moral significance of the distinction between doing and allowing. If *A* diverts the trolley towards herself, she is killing herself. I think the disparity between the good secured and the harm inflicted would have to be very significant indeed to require a person to kill herself, since suicide is an especially demanding act. But if *A* does not prevent *B* from diverting the trolley, she is letting herself die. I think that if *B* is required to divert the trolley, *A* is required to let herself be killed even if she is not required to kill herself. Again, this seems to me especially plausible once we increase the numbers to be saved in the way I suggested above. It would be wrong to try to save one's own life in a way that contributes to a lethal threat to ten or fifteen people.

McMahan has defended the permissibility of *A*'s jamming the switch by comparing it to *Surgery*:

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

It seems clearly permissible for Victim to kill Surgeon in self-defence. And *Surgery* is similar to *Trolley* in that both involve the defender's acting in a way that prevents the saving of others from something that already threatens them. And requiring restraint in both cases will entail the defenders' letting themselves die.

¹¹ In conversation. But see McMahan (2014) for a discussion of these ideas in the context of war.

But the cases are also importantly different. In *Surgey*, the harm that Surgeon will inflict upon Victim is both unjust and unjustified. Victim has a claim not to be harmed by Surgeon that he has done nothing to forfeit, and Surgeon has no objective justification for harming him. In defending himself, Victim therefore prevents himself from being harmed and wronged. Wrongs count for more in our deliberations about proportionality and lesser evil than mere harms.

By hypothesis, the harm to *A* in *Trolley* is justified. *A* is not wronged if she is killed in *Trolley*. Thus, the harm she will suffer does not count as heavily in deliberations about proportionality and lesser evil. It is this that explains why *A* may not prevent the turning of the trolley, even though Victim may prevent the saving of Surgeon's patients.

So, according to me, *B* has a duty to turn the trolley, and *A* is not permitted to prevent the turning of the trolley. Does this mean that *A* lacks a claim not to be harmed? I don't think the Requirement Thesis entails this: it's compatible with the view that *A* has a claim that is overridden. We could maintain the connection between claims and duties if we say that *B* has a *pro tanto* duty not to harm *A* that is outweighed by the duty to save the five, although this means severing the connection between duties and the standing to demand compliance.

But we might think it's nonetheless true that *A* lacks a claim in *Trolley*. A specificationist account of claims might hold that claims never give us general protection, but rather only offer limited protection in certain circumstances. So *A* has a claim not to be harmed unless, for example, harming her is a lesser evil. Rather than a murderous aggressor's having forfeited her claim not to be killed, she simply never had a claim not to be killed when she was on course to violate someone else's claim. And so on.

A third option, which is a kind of middle ground, holds that *A*'s claim persists in the face of the lesser-evil justification, but the nature of the correlative duties it confers on others can change. In *Trolley*, *B* is under no duty not to divert the trolley towards *A*. But there might be other ways in which we ought to recognize *A*'s claim. For example, there does seem to be a duty to compensate *A* (or her beneficiaries). Perhaps when making a person bear a great cost is significantly better than any other available distribution of cost, we are under no duty not to make her bear that cost. But if, after the fact, it is possible to redress the balance by compensating her, there is a duty to do so, and this duty is grounded in her right not to be harmed. This recog-

nizes that it was in some sense unfair that *A* bore all the cost, even though *B* was justified in making her bear all the cost.

We might think that this doesn't quite capture the content of the right not to be harmed. But I'm not sure that's true—after all, the point of compensation is supposed to be that it restores the recipient to their previous level of well-being. That seems like one way of reflecting the right not to be harmed. There are, of course, various possible accounts of how this compensatory duty ought to be discharged. But whichever option we choose, what matters for our purposes here is that *A*'s right not to be harmed generates some duty or duties in *Trolley*, and as long as it does this it retains the distinctive normative dimension of a claim right that Quong identifies. I find this third option quite attractive, since it seems to me less problematic than the view that *B* has both a duty not to divert the trolley and a duty to divert it, and explains why we ought to compensate *A*.

Defence in Day's End. In the three defensive variations of *Trolley* that I have discussed, it's the fact that *B* is objectively justified in diverting the trolley that undermines *A*'s claim to self-defence. In *Day's End*, in contrast, *B* is not objectively justified in flipping the light switch. Nonetheless, I think Quong's view must deny *A* a right of self-defence against *B*. This is true even if we construe *B* as an innocent threat. In other work, Quong has argued that we have agent-relative permissions to defensively harm innocent threats. In a conflict between myself and another innocent person who will kill me, I may give extra weight to my own survival. But presumably an agent-relative permission of defence obtains only if the defender has a claim not to be harmed. If, as Quong argues, *A* has no claim that *B* not harm her when his evidence is that refraining from acting would be too costly, she surely cannot, nevertheless, have an agent-relative permission to use force against *B*. She has no claim to defend, and *B* has no duty with which she can make him comply. But if this is right, it Quong's account will prohibit defence in this and other cases in which he has previously argued that it is 'plausible and coheres with many people's intuitive judgements' that defence is permissible (Quong 2012, p. 58).

I think *A* may defend herself against *B*. *B*'s mistake—reasonable or otherwise—cannot prohibit *A*'s acting in self-defence against an objectively unjustified threat. But *B*'s evidence could determine whether *B* violates *A*'s claim when he harms her. We can have a fact-

relative account of when agents have claims without committing ourselves to fact-relativity about when claims are violated. It's worth remembering that Thomson's view that *B* violates *A*'s claim is primarily motivated by a desire to sanction defensive force against *B* and other innocently threatening people. Since Thomson wants to permit defence against innocent threats, and she also thinks that defence is permissible only to prevent violations of claims, she holds that one can innocently violate a claim. But an alternative is to deny that a prospective violation of a claim is a necessary condition of permissible defence (a view that Quong and I share) (Quong 2009). If we reject this criterion, we can grant *A* a permission of defence in *Day's End* without needing to show that *B* will violate *A*'s claim if he harms her.

VII

Conclusion. I've argued that Quong is mistaken to think that the evidence available to a prospective duty-bearer determines an agent's possession of claims. Despite what Quong argues, the evidence-based view forces us to deny claims to the villagers in *Duped Soldiers*. Moreover, it's inconsistent and implausible to grant agents claims in the face of lesser-evil justifications but deny them claims in cases like *Day's End*. However, I think it's an interesting question what we ought to say about duties and claims in lesser-evil cases like *Trolley*. I've argued that if diverting the trolley is permissible, it must also be required, and *A* can have no permission of defence to prevent its being turned. But when *B* acts without objective justification, as in *Day's End*, *A* may defend herself against being harmed.¹²

Department of Philosophy
Stockholm University
Stockholm SE-10691
Sweden
helen.frowe@philosophy.su.se

¹² Drafts of this paper were presented at Trinity College Dublin, the Centre for the Study of Mind in Nature at the University of Oslo, and the University of Gothenburg. I'm grateful to the audiences for their comments. I'm especially grateful to Lars Christie for extensive discussion of these issues, and to Massimo Renzo and Seth Lazar for helpful written comments. Thanks also to Jeff McMahan and Victor Tadros for helpful discussions. And thanks to the Wallenberg Foundation for funding me during the writing of this paper.

REFERENCES

- Kagan, Shelly 1989: *The Limits of Morality*. Oxford: Oxford University Press.
- McMahan, Jeff 2014: 'Self-Defence Against Justified Threateners'. In Helen Frowe and Gerald Lang (eds.), *How We Fight: Ethics in War*. Oxford: Oxford University Press.
- Parry, Jonathan forthcoming: 'Civil Wars'. In Seth Lazar and Helen Frowe (eds.), *Oxford Handbook of the Ethics of War*. New York: Oxford University Press.
- Quong, Jonathan 2009: 'Killing in Self-Defence'. *Ethics*, 119, pp. 507–37.
- 2012: 'Liability to Defensive Harm'. *Philosophy and Public Affairs*, 40(1), pp. 45–77.
- 2015: 'Rights Against Harm'. *Proceedings of the Aristotelian Society Supplementary Volume* 89, pp. 249–66.
- Raz, Joseph 1984: 'On the Nature of Rights'. *Mind*, 93, pp. 194–214.
- Tadros, Victor 2011: *The Ends of Harm: The Moral Foundations of Criminal Law*. Oxford: Oxford University Press.
- Thomson, Judith Jarvis 1990: *The Realm of Rights*. Cambridge, MA: Harvard University Press.
- Walen, Alex, and David 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, pp. 545–71.

