

## THE LIABILITY OF MINIMALLY RESPONSIBLE THREATS

### 1. Introduction

A *minimally responsible threatener* is an agent who is responsible for imposing an objectively wrongful threat, but whose responsibility for it does not rise to the level of culpability, because they are excused for the harm they are committing. Is it morally permissible to defensively kill a minimally responsible unjust threatener when doing so is the only way to prevent the threat's victim from being killed by the threatener? And if doing so is permissible, is it because the threatener is morally liable to be defensively killed? This issue, which I address here, has remained recalcitrant in the ethics of killing in defense. Part of the purpose of addressing this issue is to determine whether both a minimally responsible unjust threatener *and* her potential victim are morally permitted to engage in lethal self-defense against one another.

According to what I call “responsibility-based accounts” of liability, minimally responsible threateners can indeed be permissibly killed. I will argue that responsibility-based accounts of liability are correct in grounding liability in a threatener's *responsibility* for a harm rather than in her *culpability* for that harm. But the standard account is nonetheless mistaken in concluding that minimally responsible lethal threateners are morally *liable to be killed*. Whereas the standard account grounds a permission to kill a minimally responsible lethal threatener in her responsibility for the harm she poses, the alternative responsibility-based account of liability which I will defend grounds a permission to kill a minimally responsible lethal threat both in that threat's responsibility for the harm she poses *and* in a lesser-evil justification for imposing an unavoidable harm on the threat rather than on the victim. Accordingly, I will call this the “hybrid responsibility-based account of liability”.

On the hybrid account, a minimally responsible unjust threatener is liable only to the amount of harm concomitant with the degree of responsibility she bears for the threat she is imposing. Anything above that amount is justified as a lesser evil, relative to the alternative of imposing the harm on her potential victim. Consequently, minimally responsible threateners can indeed be permissibly killed – but not because they are morally liable to be killed.

The arguments I will outline against the standard account and in favor of the hybrid account show that the lesser evil justifications and liability-based justifications are more closely related than has been realized. The decision to engage in self- or other-defense against minimally responsible threateners turns crucially on the lesser-evil justification – even if that decision does not affect the total number of lives that will be lost.

The hybrid account I propose has decisive advantages over standard responsibility-based accounts of liability which ground the permission to kill minimally responsible threats solely in liability. Specifically, the hybrid account ties liability more closely to responsibility, which results in a more plausible picture of *how much* defensive violence a minimally responsible threatener is morally liable to suffer. In addition, the hybrid account solves a problem with which standard responsibility-based accounts have difficulty: the hybrid account explains why it is impermissible to defensively kill *multiple* minimally responsible threats each of whom is threatening one and the same victim. (Until I discuss this, the reader should assume that the minimally responsible threateners I mention are lethally threatening just one innocent).

An upshot of the hybrid account is that, if the victim of a justified rights-infringer has an agent-relative permission to engage in self-defense against the justified attack then there can be a *moral symmetry* between minimally responsible threateners and their potential victims: both are permitted to kill the other, though for different reasons. But unlike others who have reached a similar conclusion,<sup>1</sup> the hybrid account yields this moral symmetry without abandoning the claim that minimally responsible threatener are morally liable to be harmed, or the claim that there is an agent-neutral justification for killing minimally responsible threateners.

## 2. The Responsibility-Based Account of Liability

I will begin by reviewing the responsibility-based account of liability. In doing so, I will elucidate precisely what a minimally responsible threatener is, after which I will discuss how the *standard* responsibility-based account ties liability to be killed to responsibility.

### 2.1. Responsibility, Culpability, and Minimally Responsible Threateners

As mentioned in the introduction, a minimally responsible unjust threatener is an agent who poses or contributes to an objectively wrongful threat, and bears some moral responsibility for doing so, but not to a degree that makes her culpable for that wrongful threat.<sup>2</sup> Her moral responsibility is diminished due to mitigating circumstances which fully exculpate her. But one might ask: if she is not at all culpable, then what does it mean to say that she still bears some moral responsibility for what she does? As I use the concept here, culpability *just is* a high degree of moral responsibility. Accordingly, non-culpability does not entail non-responsibility – an agent can bear some moral responsibility for a wrongful harm, despite that she is not culpable. In such a case, her moral responsibility simply has not crossed the threshold into culpability. A minimally responsible threatener occupies that middle ground: she is somewhat morally responsible, but not culpable. Here are canonical examples of minimally responsible threateners:

#### The Mistaken Resident<sup>3</sup>

A serial killer, known to kills without warning, is on the loose in a small town. Unbeknownst to anyone, the serial killer has an identical twin who stops at a random residence one night to ask for directions. Since it is a rough neighborhood and it is unusual for the resident to receive visitors this late, she answers the door armed. Seeing what appears to be the serial killer, the resident immediately shoots in what she reasonably believes to be necessary and proportionate defense of her and her family.

The harm that the mistaken resident imposes is (objectively) impermissible since shooting the killer's twin does not achieve any good and since the killer's twin has done nothing to lose her right not to be killed. Yet the mistaken resident is not culpable since she was epistemically justified in her belief that she and her family were in imminent danger; if her justified belief had not been mistaken, the harm she imposed would not have been wrongful. Though she is not culpable for the wrongful harm she commits, she retains some responsibility for it because she recognized (or was in a position to recognize) that shooting anyone – even someone who appears to be a villainous aggressor – is very morally risky act. She chose to take a moral gamble and bears some responsibility for the consequent loss of that gamble.

---

<sup>1</sup> See, in particular, (Quong, Liability to Defensive Harm, 2012).

<sup>2</sup> A threat is objectively wrong just in case it is wrong in the fact-relative sense, as opposed to the evidence-relative sense. See (Parfit, 2011, p. ch. 7).

<sup>3</sup> This is based loosely on an example by Michael Otsuka (1994, p. 91).

Here is an example of a minimally responsible threatener who, unlike the mistaken resident, imposes a harm unintentionally:

#### The Conscientious Driver<sup>4</sup>

A careful driver operating a well-maintained car is passing through a residential neighborhood on her way to the cinema. A young child playing hide-and-seek in the front yard of her parents' house unexpectedly darts out into the street, just a few meters in front of the car. The alert driver sees the child, but it is impossible to stop the car in time, even though she is driving at the speed limit. She strikes and kills the child.

Like the mistaken resident's victim, the innocent in this case has done nothing to lose her right not to be harmed – and neither does imposing the harm achieve any appreciable good. Consequently, the harm that the driver imposes on her victim is objectively wrongful; it is morally impermissible to drive on an occasion in which doing so will kill an innocent. But by driving carefully in a well-maintained vehicle, the driver was engaging in a permissible type of activity. And the driver was epistemically justified, though mistaken, in her assumption that driving on this occasion would not result in harm to an innocent. Accordingly, the conscientious driver, like the mistaken resident, is not culpable for the wrongful harm she commits. But she retains some responsibility for the wrongful harm, because she recognized (or should have recognized) that what she was doing carries a small but substantial risk of imposing a wrongful harm. Even driving safely is a risky activity; should the risk manifest, the driver is partly responsible for it by virtue of having freely chosen to undertake that risk.

There are other ways in which an individual can be minimally responsible threatener. Cognitive impairment, brainwashing, coercion, or mere youth can exculpate an individual for a wrong she commits; but so long as it is nomologically possible for her to assert control over her own actions, and to recognize that what she is doing is wrong or morally risky, she qualifies as a minimally responsible threat.<sup>5</sup> But here I will focus on epistemically justified minimally responsible unjust threateners.

The fact that minimally responsible threateners retain some responsibility for what they do distinguishes them from both *innocent* threateners and *non-responsible* threateners. An innocent threatener is someone who bears no responsibility for the harm she causes because the act does not relevantly flow from her agency. An individual whose body is thrown by a freak gust of wind onto a bystander control is an example of an innocent threat.<sup>6</sup> A non-responsible threatener is someone who causes the wrongful harm as a result of freely undertaken action, but who, unlike minimally responsible threateners, is in a position to recognize that what she is doing is morally risky. Consider the following example of a non-responsible threat:

#### The Cell Phone Caller<sup>7</sup>

---

<sup>4</sup> This is based on an example by Jeff McMahan (2009, p. 165).

<sup>5</sup> For a brief discussion of the sorts of psychological control needed to be minimally responsible, see (McMahan, 2002, p. 401) and (McMahan, 2009, p. 34).

<sup>6</sup> Judith Thomson's example of the 'falling fat man' is, of course, the prototype of this case (Thomson, 1991, pp. 287ff, 302).

<sup>7</sup> This is based on an example by McMahan (2009, p. 165). I owe the analysis of this case to him as well.

A woman's cell phone has, unbeknownst to her, been rigged so that next time she answers a call, the phone will send a signal detonating a bomb killing an innocent person. She has no more reason to believe that answering a call will detonate a bomb than you have reason to believe that your cell phone is similarly rigged.

It is objectively wrong for the cell phone caller to answer her cell phone. But doing so is subjectively permissible. Although the cell phone caller is not culpable for the objectively wrongful threat she poses – and is in that respect similar to a minimally responsible threatener – they nonetheless differ with respect to how much responsibility they bear for the threats they pose. The conscientious driver and the mistaken resident are in a position to recognize that they are undertaking a morally risky act, whereas the cell phone caller is not. Thus, should that risk actualize, the conscientious driver and the mistaken resident are at least minimally responsible for it, as a result of having knowingly undertaken a risky activity. The cell phone caller is not minimally responsible since answering cell phone calls is not in general a risky activity. Accordingly, the cell phone caller does not qualify as a minimally responsible threatener – she is not at all responsible for the harm she causes.

Now that I have discussed what a minimally responsible threatener is, I will turn to whether it is permissible to defensively attack such threateners, according to responsibility-based accounts. In doing so, first I will claim that the potential victim has decisive agent-relative reasons to defend herself against the threatener, after which I turn to whether the threatener is morally *liable* to be defensively attacked.

## 2.2. *Agent-Relative Basis for Self-Defense Against Minimally Responsible Threateners*

I believe that any responsibility-based account of lethal defense (including both the standard account as well as the hybrid account I ultimately defend) must grant that the potential victim has decisive agent-relative reasons to engage in self-defense against the minimally responsible threatener.<sup>8</sup>

Whenever one's right not to be killed is being violated by an objectively wrongful threatener who is at least minimally responsible for the threat she poses, agent-relative reasons to favor one's own life permit killing the threatener if necessary to save one's own life (provided that the impersonal harm caused by doing so is not substantially greater than the impersonal harm which would be averted). I will call this the Weak Agent-Relative Permission (or WARP) for killing in self-defense. As a consequence of WARP, the victims of the mistaken resident and the conscientious driver would all be permitted to defensively kill their respective minimally responsible threateners as a necessary means or side-effect of saving their own lives.

It is important to appreciate just how weak WARP is. First, and most obviously, a third-party would not have a permission to engage in other-defense of the potential victim either by killing the minimally responsible threatener herself or by materially aiding the potential victim. This is because WARP grounds the permission to engage in self-defense against minimally-responsible threateners in an agent-relative permission to favor one's own life. So unless the third party bears a special relation to the potential victim – one important enough to ground an agent-relative permission to favor the potential victim over the minimally responsible threatener – the third-party would be permitted to aid the potential threat only if there are agent-neutral reasons in favor of doing so.<sup>9</sup>

---

<sup>8</sup> Quong defends a thorough-going agent-relative account of self-defense in (Killing in Self-Defense, 2009).

<sup>9</sup> This point was famously made by Nancy Davis in (Abortion and Self-Defense, 1984).

(Consequently, a stranger would not be permitted to kill the conscientious driver in order to save the threatened child, but the child's parent presumably would have an agent-relative permission to do so).

Second, WARP does not claim either way whether *innocent* threateners can be permissibly killed in self-defense. It would be contentious to claim that they cannot, since it seems to make morality too demanding. And it is contentious to claim that they can be killed in self-defense, since there is seemingly no morally relevant difference between an innocent threatener and a bystander whose death is necessary to save one's own life; consequently, an agent-relative permission to kill the former would seem to apply counter-intuitively to the latter as well. Fortunately, I need not resolve this dilemma here.<sup>10</sup> WARP claims merely that a threatener's minimal responsibility is a *sufficient* condition for killing her in self-defense. This does not rule out the view that there are decisive agent-relative reasons to kill innocent threats. Nor does it over-generalize to a permission to kill bystanders, since there *is* a morally relevant difference between bystanders and minimally responsible threateners: the former, unlike the latter, bears *no* responsibility for the threat to her potential victim.

WARP is a weak principle in a third way. It does *not* claim that victims of an objectively justified rights-infringement (such as a single innocent who is killed as a foreseeable side-effect of saving twenty others) has an agent-relative permission to kill the objectively justified rights-infringer in self-defense. Though I support such a view (I return to it briefly in section 4.1) WARP itself does not imply it. Rather, WARP claims only an agent-relative permission to kill in self-defense against objectively *unjustified* minimally responsible threateners.

Quong notes that if WARP is correct, we do not need to believe that minimally responsible threateners must be *morally liable* in order to believe that their victims are morally permitted to impose lethal defensive harm on them.<sup>11</sup> While he is correct in this regard, it is still quite important to determine whether minimally responsible threateners are morally liable to be defensively killed. If they are morally liable to be defensively killed, then the minimally responsible threateners would have no agent-relative permission to engage in lethal self-defense against their potential victim's attempt at lethal self-defense. Put differently, whether there is a moral symmetry between the minimally responsible threatener and her potential victim, in that each is morally permitted to kill the other in self-defense, depends on whether the minimally responsible threatener is morally liable to be killed.

I turn to this issue – the liability of minimally responsible threatener – next. According to the standard responsibility-based account, minimally responsible threateners are indeed liable to be killed – which, I will argue, is a mistake.

### *2.3. Liability of Minimally Responsible Threats on the Responsibility-Based Account*

What does it mean to be morally liable to a harm? A person P is liable to be harmed just in case P has done something to involuntarily forfeit her right not to be harmed in that way. In addition, if P is liable to a harm, she is not *wronged* by that harm, and thus would have no justified complaint against the harmer for having been harmed. *A fortiori*, P has no right of self-defense – not even an

---

<sup>10</sup> For attempts to solve this problem in favor a permission to kill innocent threats, see (Frowe, 2008) and (Quong, Killing in Self-Defense, 2009).

<sup>11</sup> (Quong, Liability to Defensive Harm, 2012, p. 58)

agent-relative one – against the harm to which she has made herself liable. (Thus WASP does not apply to those who are liable to the harms they are facing).

According to responsibility-based accounts of liability, a minimally responsible threatener can be liable to defensive harm (and on the *standard* responsibility-based account, the harm to which the threatener is liable is a harm of death). This is because liability is grounded in responsibility; culpability is not necessary for liability. Knowingly and voluntarily engaging in an activity that imposes a small but significant risk of wrongful harms is a basis of responsibility for that wrongful harm should it manifest. And responsibility for the objectively wrongful harm grounds liability for that harm. Thus if harming a minimally responsible threatener is necessary to avert the threat she wrongfully poses, then inflicting the defensive harm is justified on the grounds that she is liable to that harm.<sup>12</sup>

On a responsibility-based account, for an agent to be minimally responsible and thus liable to defensive harm it is not enough that an agent does something objectively wrongful. Posing an objectively wrongful threat does *not* serve as a basis of liability to self-defense if the threat is posed non-responsibly. There is an added subjective element: she must recognize that what she is doing carries a moral risk. The driver's voluntary choice to foreseeably impose a risk on others serves as a basis of liability to defensive harm; the cell phone caller's voluntary choice does not foreseeably impose a risk on others, and thus does not serve as a basis of liability to defensive harm.<sup>13</sup>

The conditions of liability to defensive harm can be put more generally. On the responsibility-based account of liability a person is liable to be defensively harmed only if a) the threat, if realized, would impose an objectively wrongful harm, b) the person is at least minimally responsible for the objectively wrongful threat, c) the defensive harm is not disproportionately severe relative to the objectively wrongful threat it averts, and d) imposing the defensive harm is necessary as a means or side-effect of averting the wrongful threat.

So though minimally responsible threateners are not culpable for the objectively wrongful threats they pose, there can nonetheless be a moral asymmetry between such threats and their potential victims – an asymmetry analogous to one between *culpable* aggressors and their potential victims. In both cases, the asymmetry grounds a permission to engage in defensive violence.

So far I have claimed that on the responsibility-based account minimally responsible threateners are liable to defensive harm. But *how much* defensive harm is the threatener liable to?

On the *standard* responsibility-based account, the threatener is liable to be *killed* if necessary to save her potential victim's life. This is because, on the responsibility-based account, considerations of *fairness* partly ground a minimally responsible threatener's liability to be defensively killed. Since a minimally responsible threatener chose to put others at risk, it is fair that *she* rather than her victim should suffer the cost of the risk should it actualize – which is, in the case, death. In the words of McMahan, "Since the driver chose to impose the risk for reasons of her own, it is fair that she

---

<sup>12</sup> One might argue, however, that if the pedestrian is an adult knowingly engaging in a risky activity by having chosen to traverse near a street, then she thereby assumes a risk to herself by doing so, which affects the degree of responsibility she bears if this risk manifests. We can forestall this argument by changing the example. Suppose that the driver is crossing through a very poor village in a third world country in which cars rarely, if ever, appear. In such a case, a 'pedestrian' – that is, a villager – cannot be said to assume a risk of being struck by a car when she goes out for a walk.

<sup>13</sup> See (McMahan, *The Conditions of Liability to Preventive Attack*, 2012).

should suffer the cost rather than imposing it on the pedestrian”.<sup>14</sup> Likewise, though the mistaken resident acts reasonably given her epistemically justified beliefs, she nonetheless chooses to try killing someone knowing that the use of preventive lethal force is always risky.<sup>15</sup> (I will ultimately argue that though fairness has a role to play in determining whether lethal defense against a minimally responsible threatener is permissible, fairness does not *ground* liability but instead provides reasons to discount the disvalue of the harm that the minimally responsible threatener suffers).

Critics of the (standard) responsibility-based account have maintained that it seems strange to think that a minimally responsible threatener who has done everything that morality can fairly ask of her can be morally liable *to be killed*. It seems mistaken to claim, for example, that the conscientious driver who, in Kim Ferzan’s words, “behaves in a cautious and admirable way”, has forfeited her right to life.<sup>16</sup> David Rodin argues along similar lines, in maintaining that he is “intuitively uncomfortable” with the result that the driver is liable to defensively killed.<sup>17</sup> Quong similarly writes that “It is incredible to suppose” that the conscientious driver has “forfeited [her] rights against defensive harm and may be liable to be killed if necessary.”

The common criticism here is that the amount of responsibility needed to forfeit one’s right not to be killed cannot so radically diverge from the amount of responsibility needed to be culpable. The two should stand or fall together, or at least closer together. This criticism seems right, in that minimally responsible threats are not liable *to be killed*. However, the comparative difference in the liability between the threat and her victim grounds reason to discount the disvalue of the defensive harms imposed on the threat, resulting in a justification for killing her grounded in both the threat’s liability and in reasons to avoid a greater evil. It is to this account that I turn next.

What is at stake here is not just whether the minimally responsible threatener is liable to be killed, but whether the threatener has a permission to engage in lethal counter-defense against her potential victim’s attempt to defensively kill the threatener. For if the minimally responsible threatener is indeed liable to be killed, then WASP does not apply to her, which means that the minimally responsible threatener would have to submit herself to lethal defensive violence (either from her potential victim or from a third party).

### 3. The Hybrid Responsibility-Based Account

The hybrid responsibility-based account of liability is composed of two claims. The first is the *proportionate responsibility-based account of liability*, which limits the degree of harm to which an individual is liable to the degree of responsibility she bears for the wrongful threat she poses. The second claim is the *lesser-evil discounting view*, according to which the disvalue of the defensive harm imposed on the minimally responsible threatener who bears greater responsibility ought to be discounted.

Ultimately, an implication of this account will be that that the minimally responsible threatener’s potential victim has an *agent-neutral* permission to kill the threatener in necessary self-defense, and that the minimally responsible threatener has, in turn, an agent-relative permission to kill her potential victim in self-defense. But until the issue of moral symmetry reemerges in section 4, the sorts of permissions I discuss will be agent-neutral.

---

<sup>14</sup> (McMahan, 2009, p. 177)

<sup>15</sup> See (Otsuka, 1994, p. 91).

<sup>16</sup> (Ferzan, 2009, p. 399)

<sup>17</sup> (Rodin, 2012, p. 85)

### 3.1. The Proportionate Responsibility-Based Account

On the proportionate responsibility-based account of liability (or the “proportionate account” for short), the degree of liability that an individual bears for a wrongful threat she poses is concomitant with the degree of responsibility that she bears for posing that threat:

If P is at least minimally responsible for an objectively unjust lethal harm which she will impose on Q unless we preemptively impose a defensive harm on P, then P is liable for no more than  $n\%$  of the unjust threat she poses, where  $n$  is equal to the percent responsibility she bears for the unjust harm.

A person who is 100% responsible for a harm is fully culpable for that harm.<sup>18</sup> An individual who is only causally responsible for a harm (such as an innocent-threat) is zero percent responsible for that harm. For the sake of illustration, we can stipulate that a minimally responsible threatener is someone who is at 5% responsible for a harm. Thus anyone who is  $n\%$  responsible for a wrongful harm (where  $n > 4$ ) is liable for only  $n\%$  of the harm she wrongfully poses.

Suppose, then, that a particular minimally responsible threatener is 5% responsible for the wrongful lethal threat she poses; she is therefore liable to defensive harm *no greater* than she *would* be if she were fully responsible for a harm 5% as severe as the harm which she actually poses. For the sake of illustration, we might say that a broken arm is an example of such a harm – which means that she is actually liable to no more than she would be if she were about to culpably break someone’s arm. Presumably someone who is fully responsible for such a harm is *not* liable to be killed, even if killing her is necessary to avert the threat she poses.

So contrary to the standard account, the degree of liability that a minimally responsible threatener bears for a wrongfully posed lethal threat is smaller than the harm of death. But such an individual will still be liable to some harm. So the proportionate account, like the standard account, claims that non-culpable individuals are liable so long as they are at least minimally responsible for the threat they pose. But unlike the standard account, the proportionate account claims that the degree of liability an individual bears is no greater than the degree of responsibility she bears.<sup>19</sup>

The proportionate account is, therefore, more restrictive than the standard account. And it sets the severity of harm to which an agent is liable in a principled way – by adverting to the degree of responsibility that the agent bears for the wrongful threat she poses. So though minimally responsible threateners have forfeited their right not to suffer the proportionate harms necessary to prevent the wrongful threats for which they are at least minimally responsible, they have not forfeited their right to be *killed*. This is because they are not fully liable for the harm they actually pose. Thus *killing* a minimally responsible threatener such as the conscientious driver or the

---

<sup>18</sup> Responsibility here is not zero-sum; multiple persons can bear more than 100% responsibility for that harm. For example, each participant in a group the members of which conspire to commits a murder can be 100% responsible for that murder.

<sup>19</sup> Note that if the proportionate account is correct, a minimally responsible threatener can indeed be liable to be killed *if* the harm such a threat wrongfully poses is substantial enough. For example, if she is posing a lethal threat to *multiple* innocents, then even 5% responsibility for this threat is compatible with being liable to be killed if necessary to avert the harm she poses. I return to this issue in section 3.4.



mistaken resident is a disproportionate means to or side-effect of preventing the harms that they wrongfully pose. (That is, it violates the constraint of *narrow* proportionality).<sup>20</sup>

Why should we endorse the proportionate account, over the standard account? We can identify at least two necessary conditions for any plausible account of liability: a) an agent can be liable for a harm only if she is responsible for that harm, and b) the degree of liability than an individual bears is *limited* by the degree of responsibility she bears. Both the standard account and the proportionate account meet condition (a). But only the proportionate account meets condition (b). Recall that some theorists criticized the standard account on the grounds that it implausibly claims that an agent who bears a very low degree of responsibility for a wrongful death can nonetheless bear a very high degree of liability – she can be liable to be killed, if necessary, to avoid that wrongful death. In effect, they criticize the standard account on the grounds that it fails to meet condition (b). Of course, this doesn't mean that the degree of responsibility is *irrelevant* to the degree of liability, on the standard account – but the account still allows liability to outstrip responsibility in a way that some have found problematic, and which violates condition (b). We should endorse the proportionate account instead, because it addresses this problem by 'reigning in' liability so that it cannot outrun the agent's responsibility, the way it does on the standard account. The proportionate account sets an upper limit on how much liability an agent can bear – a limit determined in a principled way, i.e., by the degree of responsibility she bears. So, on the proportionate account, the fact that the conscientious driver (for instance) does everything that morality can fairly require of her is relevant to the function determining the degree of liability she bears. We should endorse the proportionate account over the standard account because it succeeds in meeting condition (b) whereas the standard account fails.

Though a minimally responsible threatener is not liable to be killed, I believe that doing so is still morally justified *as a lesser evil*, relative to the alternative of allowing the innocent to be killed. This is because (as I will explain in more detail) it is less bad to kill the party which is at least somewhat liable than it is to allow the party which is not liable at all to be killed, given that these are our only choices. In general, it can be permissible to harm a person, even when she has neither waived nor forfeited her right not to be harmed, if the consequences of refraining from harming her would be significantly worse from an impartial point of view. In such a case, she retains her right not to be harmed, though it is overridden or outweighed by competing moral reasons. This assumes that rights, as agent-centered restrictions, are not absolute. If the negative consequences of abiding by an agent-centered restriction are sufficiently weighty, then we have a lesser-evil justification for doing to an agent what that agent has a right that we not do. The right, in this case, is *infringed* rather than violated. But even we permissibly infringe an individual's right not to be harmed, that individual is still wronged; the reactive attitudes of resentment and anger can be morally appropriate, and we are morally required to compensate her for the harm we have inflicted. The victim might even be permitted (as I will claim) to resist the harm we are permissibly imposing on her. This is in contrast to *liability* as grounds for permissibly inflicting harms – in such a case, the individual's right is neither infringed nor violated, which is to say that she was not wronged.

The proportionate account alone does not entail that killing a minimally responsible lethal threatener is morally justified as a lesser evil. Recognizing why this is so will prove instructive. If, for example, the conscientious driver is 5% responsible for a threat of death, that means, on the proportionate account, that she is liable to however much harm it would be proportionate to inflict on her to prevent her from being fully culpable for causing a harm equivalent to 5% of the harm of death.

---

<sup>20</sup> See (McMahan, Killing in War, 2009, pp. 20-21).

Suppose, for the sake of argument, that it makes her liable to suffer defensive harm equivalent to 50% of the harm that her victim would suffer in being killed – that is, being *culpable* makes one liable to suffer a harm ten times greater than an equivalent harm committed non-culpably. Consequently, for a lesser-evil justification to permit killing the minimally responsible threatener, it would have to permit inflicting a harm equivalent to the remaining 50% of the harm she would suffer in being killed. Suppose her potential victim is 20 years old and would lose 60 years of life by being killed. Accordingly, the lesser-evil justification has to be sufficiently strong to justify inflicting on her a harm to which she is not liable that is equivalent to the loss of 30 years of good life.<sup>21</sup>

Now, a lesser-evil justification requires that the harm averted be substantially greater than the harm inflicted. Suppose we have a lesser-evil justification for intentionally killing a person who is not liable to be killed only if that were necessary to save the lives of *ten* other non-liable individuals. But if we take a multiplier of ten as required for a lesser-evil justification,<sup>22</sup> the innocent bystander would have a lesser-evil justification for killing the conscientious driver in self-defense only if she would thereby prevent herself from losing 300 years of good life.<sup>23</sup> Even if we cut in half the lesser-evil multiplier, it will still leave the lesser-evil justification with too much work to do – only in rare circumstances will reasons to avoid the greater evil, in combination with liability-based reasons, provide a permission to kill a minimally responsible threatener who is imposing a lethal wrongful harm on a single innocent. But there is an additional reason (which I turn to next) for thinking that defensively killing a minimally responsible threatener is the lesser evil relative to the alternative of allowing the threatener to kill.

### 3.2. *The Lesser-Evil Discounting View*

Deciding whether to kill a minimally responsible lethal threat involves determining whether it is better to impose a lethal harm on someone who is at least somewhat liable, or whether it is better to allow a lethal harm to befall someone who is not liable at all. In making this decision, we ought to adopt what I call the “lesser-evil discounting view”:

When determining whether there is a lesser-evil justification for imposing on a minimally responsible threatener a defensive harm greater than that to which she is liable, we ought to *partially discount* the disvalue of that harm relative to the weight of the harm that the threatener would otherwise impose on her potential victim.

So though there is an agent-centered restriction against inflicting non-liable harms, the stringency of this restriction is diminished when a) the harm is defensively imposed on a minimally responsible threatener, and b) the alternative is to allow the minimally responsible threatener to suffer a harm of equal severity.

Considerations of fairness ground the lesser-evil discounting view. Recall (from section 2.3) that proponents of the standard responsibility-based account characterized the reasons why the minimally responsible threatener is liable *to be killed* in terms of fairness. (As Quong puts it, “At the

---

<sup>21</sup> I owe this argument to Jeff McMahan.

<sup>22</sup> There is reason to believe that this multiplier is not linear with respect to the severity of the harm imposed and the severity of the harm averted. I argue for this in (Bazargan, *Varieties of Contingent Pacifism*, forthcoming).

<sup>23</sup> The harm of death consists not merely in the number of years of life deprived. But assuming as much here helps demonstrate in an especially vivid way why it is unlikely (given what has been said so far) that is a lesser-evil justification for killing a minimally responsible threat, even assuming the proportionate account of responsibility.

heart of the moral responsibility account is an appeal to distributive fairness.”<sup>24</sup> If either the threat or her victim must suffer a lethal harm, it is fairer that the one who is responsible for the predicament in the first place suffer that loss. This, it is thought, grounds the minimally responsible threatener’s liability to be killed. But I claimed that a minimally responsible threatener is liable only to the degree of harm proportionate to the degree of responsibility she bears. Fairness does not, in of itself, ground liability – *only the agent’s responsibility does that*. By attempting to ground a minimally responsible threatener’s *liability* in considerations of distributive fairness, the standard account (unlike the proportionate account) fails to match the liability of a minimally responsible threatener to her responsibility for that threat.

But fairness does indeed play a role in the ethics of defense, specifically, when it comes to determining how heavily we should weigh the out-sized, *non-liable* harms inflicted on the minimally responsible threatener. It *is* indeed less unfair for the minimally responsible threatener to bear these harms than it is for her victim to do so, because the threatener would bear some responsibility should the harm befall her potential victim, whereas her potential victim would be no responsibility should the harm befall the minimally responsible threatener.

Though the claim that it is fairer for the threatener to suffer the harm than it is for her potential victim to do so *depends* on the comparative difference between the agents’ responsibility, this claim about distributive fairness does not *ground* the minimally responsible threatener’s liability (as it does on the standard account). Instead, the claim about fairness grounds the lesser-evil discounting view, which is a not a claim about liability but about how we should distribute unavoidable harms. (The lesser-evil discounting view might seem to illicitly ‘double count’ the responsibility of the minimally responsible threatener – I address this criticism in the next section).

I have intentionally left vague the degree to which we should discount the disvalue of the harms to the minimally responsible threatener. It is clear, though, that even presuming the lesser-evil discounting view, more than a small difference in liability between a minimally responsible threatener and her potential victim is necessary for the reasons to avert the greater evil to be strong enough to permit killing the former. Suppose, for example, that the conscientious driver is imposing a threat on an adult who knew that the street she has chosen to traverse (in order to get ice-cream) is dangerous for pedestrians. Or suppose that the visitor whom the mistaken resident is threatening knows that she is the splitting image of a psychotic killer at large, and yet chooses to approach a house at the dead of night anyway. In these cases, the minimally responsible threatener’s victims bear some responsibility for the forced choice of where a lethal harm must fall. Suppose that both potential victims are at least minimally responsible for the predicament in which they find themselves, but neither is quite as responsible as those who are actually posing the threats. As a result, though the threateners are more liable than those they threaten, the difference between them isn’t great enough to override the restriction against intentional killing, even once we discount the disvalue of harms to the more liable party. This is just to say that there is no lesser evil justification for killing the threatener, *even though* it would indeed be marginally worse for that threatener’s victim to be killed. And since the threateners are merely minimally responsible for threatening the life of a single innocent, there is no liability-based justification for killing them. This is a case where we, as a third party, cannot permissibly intervene, so long as the only way to do so is by killing the threats. We have to let the harms fall where they may (though each might have an agent-relative permission to defensively kill the other, in accordance with WASP).

---

<sup>24</sup> (Quong, *Liability to Defensive Harm*, 2012, p. 56)

The difference in liability between two agents, at least one of which is minimally responsible and neither of which is culpable, must be significant in order for a lesser-evil justification to override the restriction against killing. If there is a standing moral presumption against the killing of a person, and in particular against the intentional killing of a person, it should be unsurprising if it takes more than a minuscule difference in liability for a harm to override it. But the difference in the degree of liability between a minimally responsible threatener and her non-responsible victim will always be significant (even if it is not significant enough to make the threatener liable to be *killed*) so long as the threatener is threatening her potential victim with an unjust death.

### *3.3. Combining the Proportionate Account with the Lesser-Evil Discounting View*

So far I have argued in favor of the proportionate account and the lesser-evil discounting view. When these views are combined, we have the following, which is the hybrid responsibility-based account of the permission to kill minimally responsibility threateners:

Though minimally responsible lethal threateners are not morally liable to be killed, there is still an agent-neutral permission to defensively kill them, provided that those they threaten are significantly less responsibility than they are. This is because imposing a lethal harm on someone who bears some liability is (*ceteris paribus*) the lesser evil relative to the alternative of allowing someone who bears significantly less responsibility for that harm to be killed. And this is because we ought to partially discount the disvalue of the harm imposed on the minimally responsible threatener relative to the weight of the harm that the threatener would otherwise impose on her victim.

The hybrid account claims, in effect, that though neither the minimally responsible threatener nor her potential victim is morally liable *to be killed*, it is better to defensively kill the minimally responsible threat than it is to allow her victim to be killed. Though imposing a lethal harm on the minimally responsible threatener is unjust, it would be *more* unjust to allow her to impose that harm on her potential victim. Put differently, while imposing the harm on the minimally responsible threatener *is* unjust from the first-order standpoint which considers solely whether the minimally responsible threatener has done anything to warrant that degree of harm, doing so is *not* unjust from the second-order standpoint of deciding where an unavoidable injustice should fall, due in part to the lesser-evil discounting view. The upshot is that we have a second-order lesser-evil justification to impose unavoidable lethal violence on the party who is more liable for this predicament – and that is the minimally responsible threatener.

It might seem strange that in a situation where exactly one of two persons will be killed, that there can be a lesser evil justification for choosing to kill one rather than another given that the total number killed remains the same. But more than numbers matter when assessing which of two harms counts as the lesser evil. David Rodin, among others, makes this point:

Whether action is justified on lesser evil grounds will in many cases depend on more than a simple comparison of magnitudes and probabilities of outcome harm and benefit... It is implicit in lesser evil justification that harm inflicted on an innocent bystander is a greater evil than an equivalent harm inflicted on a person liable to that harm. Similarly, harm brought about through positive agency is a greater evil than harm brought about through

negative agency, and harm intentionally inflicted plausibly counts as a greater evil than harm brought about as a foreseen but unintended side effect of justified action.<sup>25</sup>

In addition to the doing/allowing distinction and the intention/foresight distinction, the liable/innocent distinction is relevant to the lesser-evil calculation as well. All things being equal, choosing to kill someone who is morally liable to be killed is the lesser evil relative to the option of choosing to kill someone who is not morally liable to be killed. The hybrid account acknowledges the relevance of liability to the lesser-evil calculus – and by incorporating the lesser-evil discounting view, it acknowledges the relevance of fairness in deciding whether a minimally responsible threatener or her potential victim should suffer an unavoidable harm.

One might object to the hybrid view by raising the worry that discounting the disvalue of the non-liable harms inflicted on minimally responsible threats ‘double counts’ the threatener’s responsibility, since the threatener’s minimal responsibility already serves as grounds for limited liability.<sup>26</sup> Once a minimally responsible threatener has suffered the amount of harm to which she is liable, we cannot advert to her responsibility again to justify discounting the disvalue of further harms inflicted on her – or so the worry might go.

But this worry is unfounded. Regardless of who is killed – the minimally responsible threatener or her potential victim – the death is *unfair*, since neither is liable to be killed. But if the potential victim is killed, she is treated unfairly in an additional way: not only is she subjected to more harm than to which she is liable, but she is subjected to such harm despite that there is someone else – the minimally responsible threatener – who is *more* liable than she is. It is one thing for someone to suffer more harm than to which she is liable. This is unfair in the intrapersonal sense. But it is another thing for someone to suffer more harm than to which she is liable when there is someone else who is *more* liable than she is, and who is getting away ‘scot-free’. This is unfair in the interpersonal sense. To allow the less liable party to suffer the harm treats her unfairly twice over: the harm she suffers is unfair in both the intrapersonal and in an interpersonal sense. This is contrast to imposing the harm on the minimally responsible threatener – doing so treats her unfairly in only the intrapersonal sense (since she is receiving more harm than to which she is liable) and *not* the interpersonal (since there is no one else in the predicament who bears greater liability to be harmed).

The fact that it is both intrapersonally *and* interpersonally unfair for the potential victim to suffer the harm of death magnifies the badness of allowing this harm, relative to the option of shifting the harm to the minimally responsible threat. Conversely, the fact that it is *only* intrapersonally unfair and not interpersonally unfair for the minimally responsible threatener to suffer the harm of death diminishes the badness of imposing that harm, relative to the option of allowing the potential victim to suffer that harm. And this is just to describe the lesser-evil discounting view. There is, then, no illicit ‘double-counting’ of responsibility – once a person has suffered the amount of harm to which she is liable, there can indeed be remaining justification for discounting the disvalue of further harms inflicted on her, if the alternative is interpersonally unfair.

So if the hybrid account is correct, then we are permitted to kill the mistaken resident and the conscientious driver, in defense of their respective potential victims, if this is the only way to prevent these minimally responsible threateners from killing their victims.

---

<sup>25</sup> (2012, p. 97)

<sup>26</sup> Thanks to Jeff McMahan for pressing me on this point.

### 3.4. Advantage of the Hybrid Account: Residual Injustice and Multiple Minimally Responsible Threateners

A decisive advantage of the account I have developed over the standard account is its ability to explain why it is wrong to kill *multiple* minimally responsible lethal threateners who are independently threatening a single innocent. Virtually everyone – including supporters of the standard responsibility-based account – agrees that it is impermissible to kill, for example, *multiple* conscientious drivers who are independently threatening one and the same pedestrian.<sup>27</sup> However, a proponent of the standard responsibility-based account of liability is forced to say that the number of minimally responsible threateners is irrelevant; there is no point at which killing these threateners violates the proportionality constraint, since each threatener is liable to be killed. This is a problem for the standard account.

Proponents of the standard account have attempted to avoid this unpalatable conclusion by arguing that killing a minimally responsible threatener who is morally liable to be killed leaves a *residual injustice*. When multiple minimally responsible threateners are killed, these injustices aggregate in a way which eventually outweighs or overrides the victim's claim to priority. In what follows, I will explain what this means, why this argument fails, and how the hybrid account succeeds in explaining why multiple minimally responsible threateners cannot be killed.

On the standard account, a liability-based justification for harm distributes unavoidable harms in a way that comes closest to the ideally just distribution. On an ideally just distribution of the unavoidable harm, the minimally responsible threatener would be liable only to a small but significant percentage of the defensive harm, since the threatener is only minimally responsible for the harm she is imposing. (The rest of the unavoidable harm would be distributed among members of the community, precisely because, where the defensive harm is distributable, *no one* is liable for the amount over and above the small but significant portion for which the driver is liable).

But *ex hypothesi* the total harm cannot be apportioned; it must fall in its entirety either on minimally responsible threatener or her potential victim. Since there is no one else who is *more* liable than the threatener, and since the inevitable harm cannot be distributed, she ends up liable to *all* of it on the standard account. This option best approximates the ideal distribution.

Though on the standard account a minimally responsible threatener is liable to be killed, there is in McMahan's words, a "residual injustice" in killing her, since she has suffered more than her "fair share" of harm. Killing a minimally responsible threatener "leaves a moral 'remainder' in a way that killing a culpable threatener does not".<sup>28</sup> If the threatener were culpable – which is to say that if she were far more responsible for the wrongful harm she has caused – then imposing the entirety of the indivisible harm on her would *not* leave a "residual injustice" or a "moral remainder", precisely because she would be liable to suffer the entirety of the harm even if the harm were divisible.

This residual injustice supposedly provides a basis for the intuition that it is impermissible to kill *multiple* minimally responsible threats, each of which is threatening one and the same innocent. When multiple minimally responsible threats are killed, these remainders aggregate. "Given

---

<sup>27</sup> The qualification that they are acting independently is important – if they are intentionally cooperating, then they can be *complicitously* liable. See (Bazargan, *Complicitous Liability in War*, forthcoming).

<sup>28</sup> (2011, p. 554)

sufficient numbers,” McMahan says “they may eventually outweigh the defender’s claim to priority.”<sup>29</sup>

This is an ingenuous solution, but I think it is unavailable to McMahan, or any other defender of the standard responsibility-based account. The standard account does not have the resources available to articulate the “residual injustice” – at least not in a way that defeats the potential victim’s claim to priority. On the standard account, the minimally responsible threatener’s rights are neither violated nor infringed. It is true that the minimally responsible threatener is a victim of bad luck, in that she is liable to the entirety of the defensive harm partly because *it so happened* that it could not be apportioned between her and others. But the proponent of the standard account claims that the threat *is liable* to the entirety of the defensive harm necessitated by her ill-fortunate. So we cannot advert to her bad luck as grounds for thinking that the minimally responsible threatener was in any way wronged. Allowing moral luck to serve as a basis of liability *concedes* that there is nothing wrong in allowing it to serve as a basis for how we treat the agent.<sup>30</sup> Accordingly, there is no ground for subsequently claiming that treating the agent more harshly on account of her bad luck subjects her to an injustice. She might have a basis for complaint against the gods, but not against us.

On the standard account, then, there can be no residual injustice in killing minimally responsible threateners. Consequently, proponents of the account cannot explain why it is intuitively wrong to kill several minimally responsible threateners, each of whom is threatening one and the same innocent.

But the hybrid account deals with this problem easily. On this account, killing a minimally responsible threatener is permissible in part because doing so is the lesser evil. Since a lesser-evil reason for inflicting a harm is rights-infringing, the minimally responsible threatener does indeed have a basis of complaint – she was harmed in a way that exceeded the degree of harm to which she was liable. This explains why killing *several* minimally responsible threateners is wrong. If the rights-infringements are additive, there quickly comes a point where killing the minimally responsible threats simply no longer counts as the lesser evil. Indeed, I tend to think that we are not permitted to save the life of the potential victim if the only way to do so is to kill *two* minimally responsible threateners. The upshot is that the hybrid account, unlike the standard account, yields the intuitively correct result that it is impermissible to kill multiple minimally responsible threateners.

By limiting the role that moral luck plays in determining liability, the hybrid account also compares favorably to Quong’s “moral status” account of liability, according to which one is more responsible for the wrongful harms one commits if it is the result of treating another *as if* she lacks certain fundamental rights – even if this belief is epistemically justified. Thus, on this view, the mistaken resident (but not the conscientious driver) is morally liable to be killed. As with the standard view, the mistaken resident is liable to be killed since she gambles with the moral rights of another by treating her as if she lacked these rights. “Though it does not make sense to blame us when we act in these ways,” Quong says, “it does make sense to say that we may forfeit some of our rights when we take these sorts of risks.”<sup>31</sup> This means, as Quong himself notes, that a minimally responsible

---

<sup>29</sup> *Ibid.*

<sup>30</sup> That a minimally responsible threatener is liable and that her potential victim is not “is a matter of moral luck,” McMahan writes. “Yet she acted in the knowledge that bad moral luck was a possibility” – which is why her bad luck can serve as a basis of her liability” (2009, p. 177).

<sup>31</sup> (Quong, 2012, p. 66)

threatener can be liable *to be killed* even when “the available evidence may indicate there is no reasonable alternative other than using lethal force,” as is the case with the minimally responsible threatener. The moral status account, in my view, suffers from the same problem as the standard account – liability outstrips responsibility, in that it does not seem plausible to maintain that someone who has acted with positive justification in the evidence-relative sense can, by so doing, forfeit her right not to be killed. Conversely, on the hybrid account, liability is limited by one’s responsibility, which is in turn limited (among other things) by whether what one does is subjectively permissible.

One might think that the hybrid account is susceptible to an analogue of the criticism leveled against the standard account and the moral status account. Before we asked “how can someone who is doing what is morally required of her in the evidence-relative sense have thereby forfeited her right not to be killed?”. Now we might ask: “How can it be morally *permissible* to kill someone who is doing what is morally required of her in the evidence-relative sense?” But answering *this* question is not hard: on the hybrid account, infringing the rights of the epistemically justified threatener is morally permissible because failing to do so would make things go *worse* – i.e., there is a lesser-evil justification for killing the minimally responsible threatener. This account acknowledges that the minimally responsible threatener’s rights are infringed by killing her.

There is another advantage to the hybrid account. Intuitively, the minimally responsible threatener is owed compensation *ex post* for the degree of harm imposed upon her exceeding the degree to which she would be liable if the harm were distributable. In cases where defensive harm is not distributable *ex ante*, its *costs* often *are* distributable *ex post*, by way of compensation. That is, *ex post*, we are in a position to right a previous injustice – we can more justly apportion the costs of the defensive harm. But as McMahan notes, defenders of the standard account are forced to abandon the view that an individual who suffers a harm to which she is liable is not entitled to compensation for that harm.<sup>32</sup> This is a significant theoretical cost of the standard account. But on the hybrid account, the minimally responsible threatener is not liable to a harm exceeding the degree to which she would be liable if the harm were distributable. Instead, imposing the outsized totality of harm on the minimally responsible threatener is the lesser of two evils – since the alternative option is to impose the harm on the party who bears even less liability, which is an even greater miscarriage of justice. So though we are justified in imposing the harm on a minimally responsible threatener, her rights are nonetheless infringed in that the harm we are justified in imposing exceeds the degree of harm to which she is actually liable. This provides a straightforward basis for a compensatory duty – we infringed the minimally responsible threatener’s rights by treating her in a way which she was not morally liable to be treated. As a result, she is owed compensation for the degree of harm exceeding that to which she was liable.

#### **4. Conclusion: Re-emergent Moral Symmetry**

I have argued in favor of the hybrid responsibility-based account, over the standard responsibility-based account. The difference between the standard and the hybrid versions of the responsibility-based account might seem minor – but its implications are far-reaching. If the hybrid account is correct, along with what I call the “Strong Agent-Relative Permission” (SARP) then it resurrects a moral symmetry between minimally responsible threateners and their potential victims.

---

<sup>32</sup> (McMahan, 2011, p. 554)



Recall from section 2.2 that according to WARP each of us has an agent-relative permission to defensively kill an objectively wrongful threatener who is at least minimally responsible for the lethal threat she poses, so long as the balance of impersonal harms is not overwhelmingly against doing so. Here I will argue in favor of WARP's older sibling – SARP – which states the following:

If X's right not to be killed is being justifiably infringed by Y, then X has an agent-relative permission to defensively kill Y if necessary to prevent Y from killing X, unless doing would have morally catastrophic consequences.<sup>33</sup>

Whereas WARP is silent on whether the victims of justified rights-infringements can permissibly engage in self-defense, SARP claims that doing so is indeed permissible, assuming that the harm that would otherwise be averted by absorbing the cost of the rights-infringement is not catastrophic.

On SARP, *infringing* an innocent's right not to be killed serves as a basis for an agent-relative permission to engage in defensive harm. Though killing the innocent as a side-effect is justified, to ask the innocent to sacrifice herself by refraining from engaging in defensive harm is to ask too much of her (unless, of course, the good that will result is enormous). The threshold at which it is permissible to kill innocents as a side-effect of promoting the greater good is *lower* than the threshold at which such innocents are required to sacrifice themselves by refraining from engaging in self-defense (which is, in turn, lower than the threshold at which such innocents are required to sacrifice themselves by committing active suicide).

SARP does *not* entail the view, however, that those who are morally *liable* to be killed can engage in any sort of defensive harm against their attackers. Those who are liable to be killed have no permission – not even an agent-relative one – to engage in self-defense. Only those whose rights are infringed or violated have such a permission. Killing an agent who is morally liable to be killed does not infringe or violate her right not to be killed. Consequently, there are no grounds for a permission to engage in self-defense. So if the standard account were correct, then minimally responsible threateners would not have a permission – not even an agent-relative one – to engage in self-defense, even assuming SARP. But on the hybrid account, the permission to kill a lethal minimally responsible threatener is indeed compatible with an agent-relative permission for the threatener to engage in necessary self-defense.

This is because, on the hybrid account, the permission to defensively kill minimally responsible threats is grounded partly in reasons to avoid a greater evil – namely, the evil of imposing a lethal harm on a non-liable party. If those whose rights we justifiably infringe as a means or side-effect of averting a sufficiently greater (but not catastrophic) evil have an agent relative permission to defensively attack us in order to prevent their own deaths, then minimally responsible threateners can permissibly defend themselves. After all, on the hybrid account, minimally responsible threateners are not liable to be killed – rather, the permission to kill them is grounded in reasons to avert the greater evil. And unlike those who are *liable* to be killed, those who can be justifiably killed to avert a non-catastrophic greater evil are morally permitted to defend themselves against that justified attack. So if SARP is correct, minimally responsible threateners can defend themselves against the morally justified attack leveled against them.<sup>34</sup> This is contrast to the standard responsibility account; even if

---

<sup>33</sup> Quong (among others) defends this stronger agent-relative permission. See (Quong, 2012).

<sup>34</sup> This conclusion is superficially similar to Quong's, who also concludes that certain minimally responsible threateners and potential victims are each permitted to kill the other in self-defense; but on Quong's view, there is no agent-neutral

SARP is correct, minimally responsible threateners would not be permitted to defend themselves against the morally justified attack leveled against them, since on that account the minimally responsible threateners are morally liable to be defensively killed.

This means, for example, that if the conscientious driver's potential victim acts on an agent-neutral permission to engage in lethal and necessary self-defense by attempting to vaporize the driver's car, then the driver will have an agent-*relative* permission to attempt to kill her potential victim if this is necessary to forestall the victim's defensive attack; *mutatis mutandis* for the mistaken resident, or any other minimally responsible threatener who threatens a single innocent.

Of course, if the minimally responsible threatener successfully defends herself against the attempt to kill her, then the threatener will consequently commit the very harm which morally justified the attempt to kill her in the first place. One might, then, argue that a minimally responsible threat does *not* have a decisive agent-relative permission to defend herself after all, since she is responsible for the very harm which can only be averted by killing her. But we cannot appeal to her responsibility for the threat she is posing to ground the claim that she has no agent-relative permission to defend herself; we cannot do this, because she is only *minimally* responsible for the threat she poses. That she is only minimally responsible is precisely why her rights are *infringed* (albeit justifiably so) when we (or her victim) attempt to shift the inevitable harm from her victim onto the threatener. The minimally responsible threatener is, in this respect, more like an innocent bystander (such as the villagers) than she is like a culpable aggressor: the responsibility she bears for the threat is not great enough to preclude an agent-relative permission to defend herself against justified attempts to shift the threat from her victim onto her.

The upshot is that if SARP and the hybrid account are correct, there can indeed be a moral symmetry between minimally responsible threats and their potential victims: both are permitted to kill the other, though for different reasons. The potential victim of the minimally responsible threatener is permitted to kill in self-defense since killing the minimally responsible threatener is the lesser evil, whereas the minimally responsible threatener is permitted to kill in self-defense because her right not to be killed is being infringed. This stands in contrast to the standard responsibility-based account, according to which the relationship between a minimally responsible threatener and her potential victim is morally asymmetrical, in that only one of the two has a moral permission to engage in lethal self-defense.

Limiting the degree of liability that an agent bears for a harm to the degree of responsibility she bears, while acknowledging in our moral calculus the fact that it is worse to impose a harm on the less liable party, results in an account of self-defense more plausible than the standard responsibility-based account. It also reveals the central role that the lesser-evil justification plays in defense. And in combination with SARP, it yields the conclusion that though one party but not the other is liable to be defensively harmed, both are morally permitted to engage in self-defense.

---

justification for killing these minimally responsible threatener. The symmetrical defense is grounded solely in an agent-relative justification to favor one's own life. On my view, however, there is a relevant asymmetry between the two: the minimally responsible threatener but not her potential victim bears some responsibility for the predicament that they are in, which ultimately grounds (in the way that I have argued) an agent-neutral permission to kill the threatener.

Saba Bazargan  
UCSD  
2013

### Works Cited

- Bazargan, S. (forthcoming). Complicitous Liability in War. *Philosophical Studies* .
- Bazargan, S. (forthcoming). Varieties of Contingent Pacifism. In H. Frowe, & G. Lang (Eds.), *How We Fight*. Oxford University Press.
- Cody, C. A. (2008). *Morality and Political Violence*. Cambridge: Cambridge University Press.
- Davis, N. (1984). Abortion and Self-Defense. *Philosophy and Public Affairs* , 13 (3), 175-207.
- Fabre, C. (2010). Guns, Food, and Liability to Attack in War. *Ethics* , 120 (1).
- Ferzan, K. (2009). Can't Sue; Can Kill. In P. H. Robinson, K. Ferzan, & S. Gar (Eds.), *Criminal Law Conversations* (pp. 398-400). New York: Oxford University Press.
- Ferzan, K. (2013). Culpable Aggression: The Basis for Moral Liability to Defensive Killing. *Ohio State Journal of Criminal Law* .
- Ferzan, K. (2005). Justifying Self-Defense. *Law and Philosophy* , 24, 711--749.
- Frowe, H. (2008). Equating Innocent Threats and Bystanders. *Journal of Applied Philosophy* , 25 (4), 277-290.
- Frowe, H. (2011). Self-defence and the Principle of Non-Combatant Immunity. *Journal of Moral Philosophy* , 8 (4), 530-546.
- Lazar, S. (2010). The Responsibility Dilemma For Killing in War - A Review Essay. *Philosophy & Public Affairs* , 38 (2), 180-213.
- May, L. (2011). Contingent Pacifism and the Moral Risks of Participating in War. *Public Affairs Quarterly* , 25 (2), 95-112.
- McMahan, J. (2009). *Killing in War*. USA: Oxford University Press.
- McMahan, J. (2012). The Conditions of Liability to Preventive Attack. In D. K. Chatterjee (Ed.), *Gathering Threats: The Ethics of Preventive War*. Cambridge: Cambridge University Press.
- McMahan, J. (2002). *The Ethics of Killing: Problems at the Margins of Life*. Oxford; New York: Oxford University Press.
- McMahan, J. (2011). Who is Morally Liable to be Killed in War? *Analysis* , 71, 544--559.
- McPherson, L. (2004). Innocence and Responsibility in War. *Canadian Journal of Philosophy* , 34 (4).
- Otsuka, M. (1994). Killing the Innocent in Self-Defense. *Philosophy and Public Affairs* , 23 (1), 74-94.

- Overland, G. (2005). Killing Civilians. *European Journal of Philosophy*, 3.
- Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.
- Quinn, W. (1989). Actions, Intentions, and Consequences: The Doctrine of Double Effect. *Philosophy and Public Affairs*, 18 (4), 334-351.
- Quong, J. (2009). Killing in Self-Defense. *Ethics*, 119 (3), 507-537.
- Quong, J. (2012). Liability to Defensive Harm. *Philosophy and Public Affairs*, 40 (1), 45-77.
- Rodin, D. (2012). Justifying Harm. *Ethics*, 122 (1), 74-110.
- Seumas, M. (2007). Civilian Immunity, Forcing the Choice, and Collective Responsibility. In I. Primoratz (Ed.), *Civilian Immunity in War* (pp. 113-35). Oxford: Oxford University Press.
- Thomson, J. J. (1991). Self-Defense. *Philosophy & Public Affairs*, 20, 283--310.