

Killing Minimally Responsible Threats^{*}

Saba Bazargan

Department of Philosophy
UC San Diego

Abstract

Minimal responsibility threateners (MRTs) are epistemically justified but mistaken in thinking that imposing a non-negligible risk on others is permissible. On standard accounts, an MRT forfeits her right not to be defensively killed. I propose an alternative account: an MRT is liable only to the degree of harm equivalent to what she risks causing multiplied by her degree of responsibility. Harm imposed on the MRT above that amount is justified as a lesser evil, relative to allowing the MRT to kill her victim. This hybrid account, which generalizes to those who are more than minimally responsible, has considerable advantages.

1. Introduction

Suppose an individual is engaging in conduct which, she recognizes, imposes a very small but non-negligible risk of wrongful harm on another. Yet this individual is epistemically justified in thinking that undertaking the act in question is morally permissible. Suppose she is unlucky: it turns out that the small risk she knowingly imposed manifests, thereby threatening someone

^{*} I wish to thank the referees and editors at *Ethics* for invaluable criticism through the paper, as well as Jeff McMahan for comments on an earlier draft

else. The unlucky, non-negligible risk-taker is a *minimal responsibility threatener* (MRT). Is it morally permissible to kill MRTs defensively? If so, why? The answers to these questions have remained elusive in the literature on self- and other-defense. I address these questions here. In doing so, I will develop an account of defense that will apply not only to MRTs, but to those who are more than minimally responsible as well.

Here is a canonical example of an MRT:

*The Mistaken Resident*¹ – An armed and dangerous serial killer, known to kill without warning, is on the loose in a small town. Unbeknownst to anyone, he has a long-lost identical twin who happens to be passing through that same town. The twin stops at a random residence one night to ask for directions. Since the resident has been harassed in the past, 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.

The harm that the mistaken resident imposes is wrongful in the fact-relative sense² 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 for the wrongful harm she imposes since she was justified in her mistaken beliefs that she was confronted with a villainous aggressor and that the only way to save her own life was to shoot.

Suppose now that a disinterested third party can intervene just before the mistaken resident shoots the identical twin. The only options are either to

¹ See: Micahel Otsuka, "Killing the Innocent in Self-Defense," *Philosophy and Public Affairs* 23 (1994): 74-94, p. 91.

² This is Derek Parfit's terminology. He says an act "would be wrong in the fact-relative sense just when this act would be wrong in the ordinary sense if we knew all of the morally relevant facts". And an act "would be wrong in the evidence-relative sense just when this act would be wrong in the ordinary sense if we believed what the available evidence gives us decisive reasons to believe, and these beliefs were true". See: Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), 150. Some might find this distinction unhelpful, since it adverts to a third, 'ordinary' sense of wrong. Others might think that there is no purely fact-relative sense of 'wrong' -- that instead such acts, like tornadoes and diseases, are merely *unfortunate*. Given either worry, we can replace the claim that what an MRT does is wrong in the fact-relative sense with the claim that what she does has bad consequences, with no redeeming good consequences.

kill the mistaken resident, thereby saving the identical twin's life, or to do nothing, thereby allowing the mistaken resident to kill the identical twin. What should be done? This dilemma functions as a pressure point in the ethics of defense. How we respond to it forces us to clarify the moral basis for permissible defensive killing more generally.

Before I critique existing accounts and presenting my own, I will further investigate the concept of MRTs. In doing so, I contrast MRTs with *innocent* threateners and *non-responsible* threateners, and I explain the sense in which MRTs retain some moral responsibility for what they do.

2. Minimally Responsible Threateners, Innocents Threateners, and Non-Responsible Threateners

The Mistaken Resident is an example of an MRT who imposes a harm *intentionally*. But MRTs might also impose harms unintentionally:

*The Conscientious Driver*³ – A careful driver operating a car maintained as well as we can be reasonably ask of her, is passing through a residential neighborhood on her way to the cinema. Through no fault of her own, a mechanical failure causes her to lose control of her car. She strikes and kills a child who is playing in the front yard of her parents' house.

The driver's decision to drive on that occasion is wrongful in the fact-relative sense. Like the mistaken resident's victim, the child has done nothing to lose her right not to be killed, and killing her achieves no appreciable good. But the driver was justified in her mistaken belief that driving on that occasion would result in no harm. Consequently, the driver, like the mistaken resident, is not culpable for the fact-relative wrongful harm she commits.

Though the driver is not culpable, she retains some moral responsibility for the harm she commits because she was in a position to recognize that even driving safely is a risky activity. The risk is small enough that the act of driving safely is permissible on each occasion in the evidence-relative sense. But on this occasion the risk manifests, threatening a bystander. The driver

³ See Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), 165. In section 4.2 I address a variation of this case in which the victim assumes some risk of harm.

bears some moral responsibility for the manifestation of this risk by virtue of having freely chosen to undertake that risk.

The mistaken resident also bears some moral responsibility for what she does. Though she is not culpable for the wrongful harm she commits, she was in a position to recognize that shooting anyone – even someone who appears to be a villainous aggressor – is a morally risky act. Like the conscientious driver, she chose to gamble with the lives of others for her own benefit. She consequently bears some moral responsibility for the loss of that gamble.⁴

What does it mean to say that an MRT is morally responsible but not culpable? An agent is *morally responsible* for a fact-relative wrongful harm if she commits or causes that harm in a way revealing a willingness to pursue her practical commitments even when doing so non-negligibly risks violating the rights of others. This sense of moral responsibility is scalar. At the low end of the spectrum are those who commit or cause a fact-relative wrong, and whose conduct reveals a willingness to risk violating the rights of others, but whose moral deliberations are either subjectively justified or, if subjectively unjustified, are the result of coercion, deception, or an involuntary cognitive deficiency (though the actor remains aware or is in a position to be aware that what she is doing risks violating the rights of others). So an actor who actually violates the rights of others, and knowingly risked doing so, and was epistemically justified in imposing this risk, bears very little moral responsibility for what she does.

At the high end of the spectrum of moral responsibility are those who commit or cause fact-relative wrongful harms, and whose actions reveal not only a willingness to pursue practical commitments at the risk of violating the rights of others, but also reveal a failure to give due weight to the moral reasons which should have been included in her deliberations. (The standards determining the assessment of the actor's deliberations will vary with the evidence-relative wrongfulness of the harm in question). So those who cause or commit fact-relative wrongs negligently, recklessly, knowingly, or intentionally are *culpable* wrongdoers.

On this account, culpability presupposes moral responsibility, in that an individual culpably commits a fact-relative wrong *only if* her actions reveal a

⁴ See Otsuka, "Killing the Innocent in Self-Defense", 91.

willingness to non-negligibly risk violating the rights of others in pursuit of one's own ends. But moral responsibility does not presuppose culpability, because a willingness to non-negligibly risk violating the rights of others in pursuit of one's own ends can sometimes be evidence-relative permissible. This describes the mistaken resident and the conscientious driver. They are willing to put others at some risk to achieve their own certain ends. Should this risk manifest, they are morally responsible for it. But they are not culpable, since the risk they took does not reveal a failure of moral reasoning in their deliberations.

On this gloss I have collapsed culpability and moral responsibility into a single dimension: as an agent increases in moral responsibility for a fact-relative wrongful harm, she eventually crosses a threshold at which point she becomes culpable for that harm. At that point her actions do not simply reveal a willingness to put others at risk for her own ends; they also reveal that she is willing to do so even when there are evidence-relative moral reasons for her not to impose those risks. That is, they reveal a failure of moral deliberation.

The fact that MRTs retain some moral responsibility distinguishes them from both *innocent* threateners and *non-responsible* threateners. An innocent threatener bears no moral responsibility for the harm she causes because the harm does not relevantly result from her agency. An individual whose body is thrown by a tornado onto a bystander is an innocent threatener. A non-responsible threatener, on the other hand, causes a harm as a result of freely undertaken action, but the action is not a risky *type*. Moral responsibility depends not just on the features of a specific action, but on the type of action as well, since often the only epistemic access we have to whether an intended act will have specific features is by considering whether it is of a risky type. For example:

*The Cell Phone User*⁵ – A woman's cell phone has, unbeknownst to her, been rigged so that the 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.

⁵ See McMahan, *Killing in War*, 165.

It is wrong in the fact-relative sense for the cell phone user to answer her cell phone. It would be wrong in the evidence-relative sense as well if she knew that terrorists were using cell-phones in her vicinity to set off bombs. But given her situation, answering the phone is evidence-relative permissible. In this respect, she is like the conscientious driver and the mistaken resident. But whereas the driver and the resident are in a position to recognize that they are undertaking a morally risky act, the cell phone user is not.⁶ Non-culpable ignorance of the fact that an act risks imposing a harm precludes even minimal responsibility for that harm *if* the harmer is justified in believing that the act is not of a risky type. Thus the cell phone user, unlike the driver or the resident, is not even minimally responsible for the harm she imposes.

More specifically, the cell phone user, unlike MRTs, is not responsible for the death she causes because her conduct does not reveal any willingness to pursue her own ends at the cost of imposing a non-negligible risk of violating the rights of others. If I knowingly or negligently impose a *non-negligible* risk on others in furtherance of my own ends, then this reveals that I am willing to risk violating the rights of others in order to achieve my ends. Should this risk manifest, I am morally responsible for it. But imposing a ‘risk’ on others that I reasonably believe has virtually no chance of manifesting is not revealing of a willingness to put others at risk since presumably *everything* I do imposes negligible risks on others. Only when a risk is foreseeably non-negligible does imposing it reveal something about my willingness to put others at risk; so only then am I morally responsible for it.⁷ Consequently, the conscientious driver and the mistaken resident are morally responsible for the manifestation of the non-negligible risk resulting from their risky

⁶ But suppose that if the cell phone user had read the newspaper that day she would have realized that terrorists were using cell phone callers in her immediate area to trigger bombs. I might seem forced to say in this case that the caller *is* an MRT since she *is* in a position to know that answering her phone imposes a non-negligible risk. But if she is non-culpably ignorant of the fact that *refraining from reading the paper is risky*, she isn’t an MRT. (She would be an MRT if the paper was known to inform regularly its readers about common activities that are non-negligibly risk-imposing).

⁷ There might be exceptions. Suppose that each time the cell phone user answers a call she hopes that it will trigger a bomb, though she has no reason to believe that this will happen. If by chance it does, perhaps she *is* morally responsible for the harm she causes. But I’m not concerned here with threateners who *want* to do wrong.

decisions, whereas the cell phone user is not morally responsible for the manifestation of the negligible risk resulting from her decision to use the cell phone.⁸ (Again, if the cell-phone user knew or suspected that terrorists were using phones to trigger bombs, then the harm she causes would indeed be foreseeably non-negligible).

Now that I have discussed the sense in which MRTs are responsible for what they do, we are in a position to investigate extant *liability-based justifications* for killing MRTs.

3. The Liability-Based Justification for Killing MRTs

Liability-based justifications for killing MRTs have become prominent in the last decade.⁹ According to these views, MRTs such as the mistaken resident are *morally liable* to be killed. A person is morally liable to be harmed in a certain way just in case she has done something to forfeit her right not to be harmed in that way.¹⁰

⁸ Whether a risk is negligible depends on its description. The chance that driving will cause lethal harm to another is non-negligible. But the chance that driving safely in a well-maintained car will impose a lethal harm on a child playing in a front yard as a result of an unforeseen mechanical failure causing a loss of control *is* negligible. See Kimberly Ferzan “Culpable Aggression: The Basis for Moral Liability to Defensive Killing,” *Ohio State Journal of Criminal Law* (2013). I will bracket the difficult of issue of how to determine the pertinent description. But I think that a *highly specific* description of a particular accident infelicitously under-describes the various ways that even safe driving can harm others, and is consequently inapt as the proper description of the risk antecedently imposed. See Christie, C. George. “The Uneasy Place of Principle in Tort Law,” In *Philosophical Foundations of Tort Law*, edited by David G. Owen, 113-131. (Oxford University Press, 1997), at 115.

⁹ Advocates of some version of the liability-based account include Judith Thomson, Jeff McMahan, Tom Hurka, Victor Tadros, Mike Otsuka, and Helen Frowe. But these are not the only views that permit killing MRTs. Some have argued that the MRT’s victim has a decisive *agent-relative* reason to engage in lethal self-defense. See Jonathan Quong, “Liability to Defensive Harm,” *Philosophy and Public Affairs* 40, no. 1 (2012): 45-77. If this is correct, MRTs need not be morally liable to be killed in order to be permissibly killed.

¹⁰ There is a sense of ‘liability’ broader than this. Victor Tadros argues that an individual can be liable to a harm even if she has not forfeited any right. An innocent threat – such as unconscious man flung by a tornado onto someone else – can be liable to some defensive force because she would have a duty, were she able, to take on some harm to herself if necessary to prevent the

I claimed in section 2 that having knowingly and voluntarily engaged in an activity that imposes a small but significant risk on others as a side-effect of pursuing one's own ends is a basis of an MRT's moral responsibility.¹¹ Now, according to *responsibility-based* accounts of liability, an agent's moral responsibility for a fact-relative wrongful harm is enough to make her liable to proportionate harm necessary to avert the harm she is imposing.¹² Consequently, culpability is not necessary for liability on such accounts. So on responsibility-based accounts of liability, harming an MRT is permissible if it is necessary to avert the lethal threat the MRT wrongfully poses. It is permissible on the grounds that an MRT is liable to be defensively harmed.

An upshot of this view is that though MRTs are not culpable for the fact-relative wrongful threats they pose, there can nonetheless be a *moral asymmetry* between such threateners and their potential victims – an asymmetry analogous to one between *culpable* aggressors and their potential victims. In both cases the asymmetry grounds a permission for the victim or a third party – but not the aggressor – to engage in defensive violence. The chief difference between the culpable aggressor and the MRT is the degree of moral responsibility they bear for the wrongful harm, where, again, culpability is a high degree of moral responsibility for a wrongful harm. (The degree of defensive harm to which an individual is potentially liable becomes greater as her moral responsibility for that harm increases. Thus a culpable aggressor can be liable to more severe defensive harm, if necessary).¹³

Both the liability-based justification for killing MRTs and the competing hybrid justification (which I explore in the next section) presume a responsibility-based account of liability. Thus on both justifications, MRTs

harm she is imposing. See Victor Tadros, "Duty and Liability," *Utilitas* 24, no. 2 (2012): 259-277. But my focus is on liability *as rights-forfeiture*, rather than liability *as an enforceable hypothetical duty*. Those who deny that 'liability' is polysemous can re-cast my project as focusing on the role that rights-forfeiture plays in the permissibility of killing MRTs.

¹¹ For a brief discussion of the sorts of psychological control needed to be minimally responsible, see Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life*. (New York: Oxford University Press, 2002), 401 and see McMahan, *Killing in War*, 34.

¹² A terminological point: the harm that an individual is liable *to* is the harm to be imposed on her. The harm that an individual is liable *for* is the harm for which she is responsible.

¹³ I elaborate on this claim and respond to an objection in section 5.1.

are liable to some defensive harm, since MRTs bear some moral responsibility for what they do. The two accounts come apart, though, in *the severity* of the harm to which the MRT is liable. The liability-based justification assumes a *simple* responsibility-based account of liability (or “simple account” for short), according to which the party that is more responsible for a fact-relative wrongful threat is liable to suffer *that very degree of harm*, if necessary to prevent it from being imposed on her victim. Consequently, an MRT is liable to be killed if necessary to prevent her victim from being killed.

Considerations of fairness ground the simple account. Since an MRT but not her victim chose to put others at risk, it is more fair that *she* rather than her victim should suffer the cost of the risk should it actualize. On the simple account, the difference in moral responsibility between the MRT and her victim is decisive in the fair distribution of an unavoidable harm, which, in turn, grounds liability. We can summarize the simple account in this way:

The Simple Responsibility-Based Account of Liability – The party that is more morally responsible for a fact-relative wrongful threat is liable to suffer that degree of harm if necessary to prevent it from being imposed on her victim; this is because it is fairer to impose that harm on the more responsible party than it is to allow her to impose that harm on the less responsible party.

If the simple account is correct, MRTs such as the conscientious driver or the mistaken resident are liable to be killed, considering that the alternative would be to allow them to kill their non-responsible victim. And this is just to state the liability-based justification for killing MRTs. But the simple account (and thus the liability-based justification) is mistaken – or so I will argue. First, though, I will present my alternative account.

4. The Hybrid Justification for Killing MRTs

The hybrid justification for the killing of MRTs is composed of two principles. The first is the *complex* (as opposed to the *simple*) responsibility-based account of liability. The second is the *lesser-evil discounting view*. I will explain and defend each principle; then I will explore their implications.

4.1. *The Complex Responsibility-Based Account of Liability*

On the complex responsibility-based account of liability (i.e., the “complex account”), the degree of harm to which an individual is liable is the harm that she threatens to cause multiplied by the percentage degree of her moral responsibility:

The Complex Account of Liability – If P is at least minimally responsible for an objectively unjust harm which she will impose on Q unless we preemptively harm P, then P is liable for no more than n% of the unjust harm for which she is responsible, where n is equal to the percent moral responsibility she bears for that unjust harm.

Suppose that we can quantify the degree of moral responsibility that an individual bears for a wrongful harm. For present purposes, if an individual is 100% responsible for a harm is then she is fully culpable for it. An individual who is only causally responsible for a harm – such as an innocent-threat – is zero percent responsible for it. For the sake of illustration, we can stipulate that an MRT is someone who is between 5% and 10% responsible for a harm. Thus anyone who is n% responsible for a wrongful harm (where $n \geq 5$) is liable for only n% of the harm she wrongfully poses.

There are three ways in which this claim is idealized. First, the claim that MRTs are, specifically, between 5% and 10% responsible for a harm is illustrative in the sense that we might determine that the lines should be drawn elsewhere. Second, I am assuming that harms can be measured on a single dimension yielding an interval measure of their moral significance. Third, I am assuming that moral responsibility can be quantified. While the claim that an individual is, for example, 16% as opposed to 17% responsible for a harm constitutes a false precision, I think it is plausible to maintain that moral responsibility can be *coarsely* quantified. In such a case, the intervals might be five percent or ten percent. A consequence of coarsely quantifying moral responsibility is that a *slight* difference (for example, a difference of one or two percent) will not determine whether an individual qualifies as an MRT. This is important, since whether an individual so qualifies makes a significant difference in how she can be treated; consequently, whether she qualifies as an MRT should not depend on

drawing sharp distinctions between very slightly differing degrees of moral responsibility.

With these idealizations in mind, suppose, then, that a particular MRT is 5% responsible for the 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 she actually threatens. For illustrative purposes, we might say that a broken wrist exemplifies such a harm. Consequently the MRT is actually liable to no more than she *would* be if she were fully culpable for *attempting* to break someone's wrist (i.e., if she were not merely minimally responsible).¹⁴ Presumably someone fully culpable for such a harm is *not* liable to be killed. So contrary to the simple account, the degree of harm to which an MRT is liable is less than death. (Of course, if the MRT will die one day later in any case, or has a life barely worth living, the defensive harm imposed on might *not* be excessive relative to the risk she imposes -- but I am assuming that killing the MRT deprives her of a life well worth living).

The complex account limits the degree of harm to which an individual is liable in a principled way – by advertent to the degree of moral responsibility that the agent bears for the wrongful threat she poses. Consequently, the account is more restrictive than the simple account. Though MRTs have forfeited their right not to suffer the proportionate harms necessary to prevent the wrongful threats they are posing, they have not forfeited their right to be *killed*. This is because they are not responsible enough for the threats they pose. Thus killing an MRT such as the conscientious driver or the mistaken resident is disproportionate. (Specifically, it violates *narrow* proportionality).¹⁵

The complex account differs not only from the simple account, but from Jonathan Quong's "moral status" account of liability as well.¹⁶ On the moral

¹⁴ If it turns out that the moral evaluation of harms or degrees of responsibility do not admit of interval orderings, the complex account can be re-characterized as claiming that lower responsibility results in liability to less severe harms; such a claim does not require quantifying responsibility or harms.

¹⁵ For an explanation of narrow proportionality, see *Ibid.*, 20-21.

¹⁶ Quong, "Liability to Defensive Harm"

status account, an agent is more responsible for the harms she commits when it is the result of treating another *as if* she lacks certain fundamental rights – even if this belief is epistemically justified. This describes the mistaken resident’s but not the conscientious driver’s conduct, insofar as the former but not the latter intentionally harms her victim. So on Quong’s view, an MRT such as the mistaken resident is liable to be killed. This is in contrast to the complex account, which says that the mistaken resident is liable only to harm significantly less severe than death.

With respect to the conscientious driver, Ferzan¹⁷ and Rodin¹⁸ (among others) argue that it’s implausible to think that she has forfeited her right not to be killed after having done everything that morality can fairly ask of her. So these critics might favor the complex account over the simple account since the former is less draconian; unlike the simple account, the complex account does not imply that the conscientious driver has forfeited her right not to be killed. But I suspect that they would not be wholly satisfied since, on the complex account, MRTs can still be liable to be killed if they are responsible for a catastrophic harm to *many* innocents. For example, suppose J is paid to transport volatile materials by freeway. Though she takes all reasonable precautions, there is still a small but significant chance of a large explosion. It occurs, threatening the lives of 20 others. Suppose J is only 5% responsible for this harm. On the complex account J is liable to a degree of defensive harm no greater than the amount to which she would be liable if she were fully responsible for the death of *one* – which is 5% of 20 deaths. So J is liable to be killed, not only on the simple account, but on the complex account as well.¹⁹ Though I find this consequence wholly plausible, others (such as Rodin and Ferzan) might demur, (though of course they could maintain that though J is not liable there is a lesser-evil justification for killing J).

¹⁷ Kimberly Ferzan, “Can’t Sue; Can Kill,” In *Criminal Law Conversations*, edited by Paul H. Robinson, Kimberly Ferzan and Stephen Gar, 398-400. (New York: Oxford University Press, 2009), at 399.

¹⁸ David Rodin, “Justifying Harm,” *Ethics* 122, no. 1 (2012): 74-110, at 85

¹⁹ I am assuming that harms to separate individuals can be aggregated, and in so doing I am using simple summation as the relevant function. I do not mean to rule out competing aggregation functions.

Before turning to the second part of the hybrid justification, I will address two criticisms of the complex account.

Criticism One

The complex account has implications for threateners who are *more* than minimally responsible. It states that an individual is liable only *for* that which she is responsible, and that the amount of harm *to* which she is liable is limited by her moral responsibility. This seems to suggest that *none* of twenty *fully responsible* threateners who conspire to push a boulder off a cliff onto an innocent would be liable to be killed defensively, if less than twenty of them are enough to do the job. Killing any of them would seem to violate proportionality, because contributing 1/20th of the force sufficient for pushing the boulder off the cliff cannot make one liable to be killed, and because no one of the responsible threateners is a but-for cause of the murder. It seems that, contrary to the complex account, proportionality should be derived from the sum of the harms inflicted, rather than from what any single contributor does.

But the complex account is compatible with this view. On an account of *complicitous* liability, each contributor is responsible for more than what she alone causes; she is also responsible for the contributions made by her co-conspirators as well. Responsibility in such cases is not zero-sum – each can be fully responsible for the jointly committed harm.²⁰ This is because cooperation in furtherance of a joint goal inculcates each participant, making each responsible for what the other does.²¹ This is reflected in Anglo-American criminal law: each participant would be treated as a co-principal, equally responsible for murder.

But suppose there is no conspiracy. Instead, multiple persons act in ways that, if each acted alone, would pose no significant harm, but which together foreseeably but intentionally combines with similar acts committed by others resulting in a substantial harm. If each contribution is a but-for or actually

²⁰ See Michael J. Zimmerman, “Intervening Agents and Moral Responsibility,” *The Philosophical Quarterly* 35, no. 141 (1985): 347-358, at 355 and Christopher Kutz, *Complicity: Ethics and Law for a Collective Age* (Cambridge University Press, 2000).

²¹ I’ve argued for this view elsewhere. See Saba Bazargan, “Complicitous Liability in War,” *Philosophical Studies*, (2012).

sufficient cause of that harm, and each knows or is in a position to know this, then the contributor is morally liable to be killed on the complex account – which strikes me as intuitively correct. If, alternatively, the jointly caused harm is causally over-determined in that each individual contribution is marginal, then it is not obvious that the complex account yields the wrong conclusion when it says that each contributor is responsible for no more than her own contribution.

Criticism Two

Quong criticizes what he calls the “variable culpability account” which says that the degree to which a person is liable to defensive harm varies with her culpability.²² His criticism might seem to generalize to the complex account. In Quong’s example, Carl unjustly threatens to break Albert’s legs unless Albert cuts off Betty’s legs. The variable culpability account says that since Albert is not fully culpable for his attack on Betty, the degree of defensive harm to which he is liable is lower than it would be if he were fully culpable. Suppose that someone fully culpable for what Albert is doing would be liable to be killed. Since Albert is less than fully culpable, it seems, then, that on the variable culpability account he cannot be liable to be killed even if killing him is necessary to save Betty’s life. This is an unintuitive result. And since, on the complex account, the harm to which one is liable varies with the degree of one’s moral responsibility and thus with culpability, Quong’s argument seems to apply to the complex account as well.

But Quong’s argument is persuasive only if we mistakenly suppose that the most that one could justifiably do to a *fully* culpable person in this case is to kill her. This is implausible, since there are fates worse than death to which a fully culpable person -- such as Carl -- could be liable. So there is room to maintain that even though Albert bears enough culpability to make him liable to be killed, he might not bear enough to make him liable to be *tortured for a day and then killed*, even if that were the only way to prevent him from cutting off Betty’s legs. This is precisely because Albert is not fully culpable, *ex hypothesi*. If Albert *were* fully culpable, however (if he were as culpable as Carl, for example) then he might indeed be liable to be tortured

²² Quong, “Liability to Defensive Harm”, 50-51

for a day and then killed, if that were the only way for Betty to save her legs.

The upshot is that the variable culpability account does not yield the unintuitive conclusion that people who are less than fully culpable are not liable to be killed. Rather, it yields the conclusion that those who are less than fully culpable will, *ceteris paribus*, be liable to less severe harms than those who are fully culpable. And this is compatible with the view that they are *both* liable to be killed.

4.2. *The Lesser-Evil Discounting View*

Recall that on my view, though MRTs are not liable to be killed on the complex account, killing them is still morally justified *as a lesser evil*, relative to the alternative of allowing the MRT's victim to be killed. 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. In such a case, she retains her right not to be harmed, though it is outweighed by competing moral reasons. (This assumes that rights are not absolute). If the negative consequences of abiding by a constraint 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.²³

Though neither the MRT nor her victim is liable to be killed, killing the MRT is the lesser evil since the outcome in which the party who is at least somewhat morally responsible is killed is less bad than the outcome in which the party who is not morally responsible at all is killed. However, the complex account does not by itself yield the view that there is a lesser-evil justification for killing the MRT. Recognizing why this is so will prove instructive.

²³ A right is infringed (rather than violated) when there is an all-things-considered moral justification for wronging the victim. For example, by switching a trolley track away from the five and toward a single innocent, we infringe but do not violate the rights of the single innocent.

Suppose, for example, that the conscientious driver is 5% responsible for the threat of death she imposes. On the complex account, she is liable to the amount of harm it would be permissible to inflict on her to prevent her from *fully culpably* imposing a harm equivalent to 5% of the harm of death. Suppose, for the sake of argument, that being fully culpable for a harm h makes one liable to suffer a harm ten times greater than the harm to which one would be liable if one had imposed h non-culpably. So the driver is liable to suffer a defensive harm equivalent to ten times 5% of the harm that his victim would suffer. That is, the driver is liable to suffer a defensive harm equivalent to 50% of the harm that her victim would suffer (since 5% multiplied by ten is 50%). Consequently, for a lesser-evil justification to permit killing the MRT, 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 the MRT a harm to which she is not liable that is equivalent to the loss of 30 years of life. That is a lot of harm that has to be justified on the ground that it is the lesser evil.

To make matters worse, a lesser-evil justification requires that the harm averted be substantially greater than the harm inflicted, due to the moral relevance of the doing/allowing distinction. For example, we might say that 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,²⁴ we would have a lesser-evil justification for killing the MRT only if doing so would prevent the victim from losing 300 years of good life (since 30 years times 10 is 300 years).²⁵ Even if we cut the multiplier in half, it will still leave the lesser-evil justification with too much work to do.

²⁴ 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. See: Saba Bazargan, "Varieties of Contingent Pacifism," In *How We Fight*, edited by Helen Frowe and Gerald Lang. Oxford University Press, forthcoming.

²⁵ Of course, the harm of death consists not merely in the number of years of life deprived. But assuming as much helps demonstrate in an especially vivid way why it is unlikely (given what has been said so far) that there is a lesser-evil justification for killing an MRT.

In spite of all this, killing the MRT is still the lesser evil. This is because in weighing the alternative to killing the MRT we ought to adopt this principle:

The Lesser-Evil Discounting View – When determining whether there is a lesser-evil justification for imposing on an MRT a defensive harm greater than that to which she is liable, we ought to *discount* the disvalue of that harm relative to the weight of the harm that the threatener would otherwise impose on her potential victim.

The lesser-evil discounting view diminishes the amount of wrongful harm which must be averted for there to be a lesser-evil justification for killing an MRT. It does this by claiming that the harm imposed on an MRT – the harm exceeding that to which she is liable – should be discounted.

Considerations of distributive justice explain why we should do this discounting. We have to weigh these two exclusive options against each other: a) the badness of imposing harms on an MRT more severe than that to which she is liable, and b) allowing those harms to befall the MRT's non-responsible victim. It is better for the MRT to bear those harms than it is for her victim to do so, because the threatener bears some moral responsibility for the predicament they both face – that is, the threatener's action is what imposes the forced choice of having to decide whether a harm should befall her or her potential victim. Insofar as one of the two parties has to bear a cost, and the threatener bears some moral responsibility for the predicament whereas her victim does not, it is fairer for the threatener to be killed than it is for her victim to be killed.

This reasoning should sound familiar – recall (from section 3) that on the simple account considerations of distributive justice ground the view that an MRT is *liable* to be killed. Similar reasoning grounds the lesser-evil discounting view: 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.

Though the lesser-evil discounting view and the simple account are grounded in similar reasoning, they represent competing views of how we ought to incorporate considerations of distributive justice into an account of defense. Whereas the simple account incorporates considerations of distributive justice directly into its account of liability, the hybrid account incorporates it into

the lesser-evil calculation. This seemingly abstruse difference yields substantial theoretical benefits, and has significant implications, which I explore in the sections 5.1 and 5.2.

It is important to note that a non-negligible difference in moral responsibility between an MRT and her potential victim is necessary for there to be a lesser-evil justification for killing the MRT. Suppose, for example, that the conscientious driver is posing a threat to an adult who knew that the street she has chosen to traverse (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. Call these “reckless victims”.

Reckless victims have assumed a risk of substantial harm. This diminishes the badness of harming a reckless victim when doing so is necessary to achieve some appreciable good. To determine whether we should kill the MRT or allow the MRT to kill the reckless victim, we have to compare the degree of moral responsibility that the reckless victim bears for the risk, with the degree of moral responsibility that the MRT bears for imposing the fact-relative wrongful threat. Suppose it turns out that the MRT bears slightly more moral responsibility than the reckless victim does. In such a case there is no lesser evil justification for killing the MRT, *even though* it would indeed be marginally worse for the reckless victim to be killed. This is because the slight difference in moral responsibility between the MRT and the reckless victim cannot override the restriction against intentional killing, even once we discount the disvalue of the harms imposed on the more responsible party. So in this case, not only is the MRT not liable to be killed, but she also cannot be permissibly killed at the bar of the lesser evil justification. This is a case where a third party cannot permissibly intervene if the only way to do so is by killing the MRT. We have to let the harms fall where they may.

4.3. Combining the Complex Account with the Lesser-Evil Discounting View

So far I have described and defended the complex account and the lesser-evil discounting view. These two combined yield:

The Hybrid Justification for Killing MRTs – Though MRTs are not morally liable to be killed, there is an agent-neutral permission to kill them defensively. This is because imposing a lethal harm on someone who bears some moral responsibility for an unjust threat is (*ceteris paribus*) the lesser evil relative to the alternative of allowing the MRT to kill her non-responsible victim. The lesser evil is substantial enough to provide a justification for killing the MRT since the disvalue of the harm imposed on the MRT is discounted relative to the weight of the harm that the MRT would otherwise impose on her victim.

Though killing an MRT *is* unjust from the first-order standpoint which considers solely whether the MRT 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. This is due in part to the lesser-evil discounting view. Consequently, we have a second-order lesser-evil justification to impose unavoidable lethal harm on the party who is more responsible for this predicament – and that is the MRT.

It might seem strange that where exactly one of two persons will be killed, 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 the lesser evil. David Rodin, among others, makes this point:

The value of states of affairs relevant to lesser evil justification depends crucially on how harms and benefits come into being, and how they are related to human agency. It is implicit in a 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.²⁶

Like the doing/allowing distinction and the intention/foresight distinction, the liable/non-liable distinction is relevant to the lesser-evil calculation. *Ceteris paribus*, killing someone is who is morally liable to sub-lethal harm is

²⁶ Rodin, “Justifying Harm,” 97

the lesser evil relative to the option of killing someone who is not morally liable to be harmed at all. 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 distributive justice in determining whether an MRT or her potential victim should suffer the unavoidable harm.

Though both the simple account (which I rejected) and the lesser-evil discounting view (which I endorse) are grounded in considerations of distributive justice, the accounts differ dramatically in *what role* distributive justice plays in self- and other defense. Proponents of the simple account treat self-defense as a distribution-problem: how do we distribute an inevitable harm for which one party bears some moral responsibility? On this construal of the problem, we impose the harm on the more responsible party, since that is fairer than the alternative. The claim that it is fairer for the threatener to suffer the harm than it is for her potential victim to do so is what grounds the threatener's liability. The simple account folds considerations of fairness into the account of liability.

But on the hybrid account, liability is not the output of a distributive problem. Considerations of distributive justice play no role in determining who is liable or to how much harm she is liable. Rather, on this view, liability is determined by a) the degree of moral responsibility that an individual bears for a fact-relative unjust harm, and b) the severity of that harm. The problem of distribution still exists, but it is dealt with by a lesser-evil calculus; comparative differences in moral responsibility constitute one of the many factors included in that calculation which ultimately determines who should be killed.

One might raise the following objection to this view. Once an MRT has suffered the amount of harm to which she is liable, it seems we cannot invoke her moral responsibility again to justify discounting the disvalue of further harms inflicted on her – this illicitly 'double counts' her moral responsibility. But this worry is unfounded. Regardless of whether the MRT or her potential victim is killed, 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 even though there is someone else –

the MRT – who is *more* responsible than she is. It is one thing to suffer more harm than to which you are liable. This is unfair in the absolute sense. But it is another thing to suffer more harm than to which you are liable when there is someone else who is *more* responsible than you are, and who is getting away ‘scot-free’. This is unfair in the comparative sense. To allow the less responsible party to suffer the harm treats her unfairly twice over: the harm she suffers is unfair in both the absolute and the comparative sense. This is in contrast to imposing the harm on the MRT – doing so treats her unfairly in only the absolute sense.

The fact that it is *only* absolutely (and not comparatively) unfair for the MRT to suffer the harm of death diminishes the badness of imposing that harm on her 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 moral 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 comparatively unfair.

If the hybrid account is correct, then we are permitted (and in some cases required) to kill the mistaken resident and the conscientious driver, in defense of their respective potential victims. But consider a variant of Conscientious Driver:

The Paramedic – A paramedic is driving an ambulance alertly to the site of a reported heart-attack victim. The ambulance, though maintained as well as can be reasonably asked, suffers a mechanical failure which causes the paramedic to lose control of the vehicle through no fault of her own. The paramedic consequently strikes and kills a child who is playing in the front yard of her parents’ house.

Like the conscientious driver, the paramedic is not to blame for the fact-relative wrong of killing the child. Yet, like the conscientious driver, she retains some moral responsibility for the wrongful harm she commits. If killing the paramedic results in the death of two persons – the paramedic herself and the heart-attack victim which the paramedic would have otherwise reached, then it is presumably impermissible to kill the paramedic on either the liability-based or the hybrid account.

But suppose that killing the paramedic will not endanger anyone else. (The apparent heart-attack victim is merely suffering from heartburn). Some argue that it is a counterintuitive consequence of responsibility-based accounts of liability that it permits killing the paramedic in such a case.²⁷ But responsibility-based accounts need not yield such a view. We as citizens have reasons to assume some risk of fact-relative unjust harm from ambulance-drivers since a) we are made better off by living in a society where such emergency services are prevalent, and b) we assume the attendant risk of harm by choosing, if only hypothetically, to be made better off in such a way.²⁸ If the potential victim assumed some substantial risk of harm, the difference in the degree of moral responsibility between the paramedic and her potential victim for the lethal harm will not be great enough to overcome the constraint against intentional killing. As stated earlier, the difference in moral responsibility between two agents, at least one of which is minimally responsible and neither of which is culpable, must be non-negligible in order for a lesser-evil justification to override the restriction against killing.

So there is no agent-neutral permission to kill the paramedic on the hybrid account since she is engaging in an activity of a type that serves a positive social function, for the sake of which we all assume some risk of harm. In contrast, driving in general does not serve a positive social function great enough to outweigh the associated harms.²⁹ Though the specific act of driving the ambulance *on that occasion* has no substantial positive benefit, shifting the threat that the paramedic non-culpably but responsibly poses back onto her would violate a hypothetical agreement we have made to accept the risks of non-culpable ambulance driving.³⁰

²⁷ Quong, "Liability to Defensive Harm," 57

²⁸ For more on how hypothetically chosen regimes can affect responsibility, see Alan Wertheimer, "Intoxicated Consent to Sexual Relations," *Law and Philosophy*, (2001): 373-401.

²⁹ See Julia Meaton and David Morrice, "The Ethics and Politics of Private Automobile Use," *Environmental Ethics* 18 (1996): 39-54 and Doug Husak, "Vehicles and Crashes: Why is this Moral Issue Overlooked?" *Social Theory and Practice* 30, no. 3 (2004): 351-370.

³⁰ But suppose we restrict the function of driving to enabling people to undertake activities that in the final calculus serves a net positive social function. Even in this case, I suspect that a conscientious driver would be liable, precisely because there is an alternative hypothetical scheme -- *viz.*, mass transportation -- that serves the same function with fewer risks and less

5. Advantages and Implications of the Hybrid Justification

So far I have described the hybrid justification for killing MRTs. In what follows, I present reasons for adopting it, and explore its implications.

5.1. *Residual Injustice, Multiple MRTs, and Redress*

A decisive advantage of the hybrid justification (over the liability justification) is its ability to explain why it is wrong to kill *multiple* MRTs who are independently threatening a single innocent. (I have in mind a case where each MRT is posing on her own a threat of a distinct harm. In each case the threat that the MRT is posing is necessary and sufficient for the occurrence of the harm, provided that the MRT isn't stopped). Even if there is intuitive disagreement as to whether MRTs can be permissibly killed, virtually everyone – including supporters of the liability-based justification – agrees that it is impermissible to kill, for example, multiple conscientious drivers independently threatening one and the same pedestrian with death. However, a proponent of the liability-based justification seems forced to say that the number of MRTs is irrelevant – there is no point at which killing these threateners violates proportionality, since each is liable to be killed. This is a big problem for the liability-based justification.

Proponents of the liability-based justification have attempted to avoid this problem by arguing in the following way. Suppose only one MRT is threatening a victim. The sum total of the harm that MRT imposes cannot be lessened but it can be divided up and distributed. Ideally, the an MRT would be liable only to a small but significant percentage of the defensive harm, since she is only minimally responsible for the harm she is imposing. The rest of the unavoidable harm would be distributed among everyone else in the world, precisely because, where the defensive harm is distributable, *no one* is liable for the amount over and above the small but significant portion

harm. Such an alternative is presumably unavailable for paramedics to reach destinations in an emergency.

for which the MRT is liable. But since the inevitable harm cannot be apportioned, and since there is no one else who is *more* responsible than the threatener, she ends up liable to *all* of it on the simple responsibility-based account. Imposing the harm on her best approximates the ideal distribution. Still, proponents of this account admit that killing the MRT leaves, in McMahan's words, a "residual injustice" since by killing her she suffers more than her "fair share" of harm – i.e., the amount of harm that should be imposed on her if the harm were divisible.³¹ (If the threatener were culpable then imposing the entirety of the indivisible harm on her would *not* leave a "residual injustice" 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* MRTs, each of whom is threatening one and the same innocent. When multiple MRTs are killed, these residual injustices aggregate. "Given sufficient numbers," McMahan says "they may eventually outweigh the defender's claim to priority."³²

This is an ingenuous solution, but it is unavailable to McMahan or any other proponent of the liability-based account. The account has painted itself into a corner. It does not have the resources to articulate the "residual injustice" – at least not in a way that defeats the potential victim's claim to priority. On the liability-based account, the MRT's rights are neither violated nor infringed when she is killed. It is true that the MRT 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 on the liability-based account the MRT *is liable* to the entirety of the defensive harm necessitated by her ill-fortune. So we cannot refer to her bad luck as grounds for thinking that the MRT 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. Accordingly, there is no ground for subsequently claiming that treating the agent more harshly on account of her bad luck treats her unjustly. There is, then, no residual

³¹ Jeff McMahan, "Who is Morally Liable to be Killed in War?" *Analysis* 71 (2011): 544-559, at 554

³² *Ibid.*

injustice. The MRT might have a basis for complaint against the gods, but not against us.

Proponents of the liability-based justification cannot explain, then, why it is intuitively wrong to kill several MRTs, each of whom is threatening one and the same innocent.³³ But the hybrid account deals with this problem easily. On this account, killing an MRT is permissible in part because doing so is the lesser evil. Since a lesser-evil reason for inflicting a harm is rights-infringing, the MRT does indeed have a basis of complaint – she is harmed in a way exceeding the degree of harm to which she was liable. This explains why killing *several* MRTs is wrong. If the rights-infringements are additive, there quickly comes a point where killing the MRTs simply no longer counts as the lesser evil. The upshot is that the hybrid account, unlike the simple account, yields the intuitively correct result that it is impermissible to kill multiple MRTs.³⁴

There is another advantage to the hybrid justification. Intuitively, an MRT is owed compensation for the amount of harm imposed upon her exceeding the degree to which she would be liable if the harm were distributable. We are in a position, *ex post*, to right a previous injustice: we can more justly apportion the costs of the defensive harm. But if the MRT (or more aptly, her estate) is owed compensation, and if the MRT is liable *to be killed*, then we are forced to abandon the principle that an individual cannot be entitled to compensation for a harm to which she is liable.³⁵ Abandoning this

³³ Note that this problem applies also to a view similar to but importantly different from the simple responsibility-based account of liability (pointed out to me by an editor at *Ethics*). It might be thought that when I drive I waive my right to be killed in defense against any lethal threat I pose while driving. But again, on this view, it will turn out permissible to kill multiple MRTs.

³⁴ One might attempt to salvage the liability-based justification by noting that killing MRTs would harm third parties, such as their families and friends, who (unlike the threateners) are not liable to be harmed. If many MRTs are killed, then the harms to the third parties aggregate, ultimately outweighing the potential victim's claim to priority. But this argument reaches the right conclusion for the wrong reasons: it is intuitively impermissible to kill many MRTs even when they have no families or friends.

³⁵ McMahan admits that his view has this consequence. See McMahan, "Who is Morally Liable to be Killed in War?" 554.

plausible principle represents a significant theoretical cost of adopting the liability-based justification.

But on the hybrid justification, the MRT is not liable to a harm exceeding the degree to which she would be liable if the harm were distributable. Instead, killing the MRT is the lesser of two evils. So though we are justified in imposing the harm on the MRT, her rights are nonetheless infringed in that the harm we are justified in imposing exceeds the amount of harm to which she is actually liable. *This provides a straightforward basis for a compensatory duty* – we infringed the MRT’s rights by treating her in a way which she was not morally liable to be treated. As a result, she (or her estate) is owed compensation for the degree of harm exceeding that to which she was liable.

In summary, then, there are two advantages to the hybrid justification over the liability-based justification – it explains why multiple MRTs cannot be permissible killed, and it grounds the view that an MRT is owed compensation for the amount of harm imposed upon her exceeding the degree to which she would be liable if the harm were distributable.

5.2. *Re-emergent Moral Quasi-Symmetry*

An interesting implication – which might be welcomed or not depending on one’s intuitions – is that the hybrid justification supports a qualified moral symmetry between MRTs and their potential victims. Arguably, if X’s right not to be killed is being justifiably infringed by Y, then X has an agent-relative permission to kill Y defensively if necessary to prevent Y from killing X (unless doing so would have morally catastrophic consequences).³⁶ On this view, which I call “the agent-relative principle”, an all-things-considered morally justified lethal attack against an agent is compatible with a

³⁶ Nancy Davis (among others) defends this agent-relative permission. See Nancy Davis, “Abortion and Self-Defense,” *Philosophy and Public Affairs* 13, no. 3 (1984): 175-207. McMahan, at one point, also held this view (see Jeff McMahan, “The Basis of Moral Liability to Defensive Killing,” *Philosophical Issues* 15 (2005): 386-405, at 388). Quong holds it as well (see Jonathan Quong, “Killing in Self-Defense,” *Ethics* 119, no. 3 (2009): 507-537) as does Uwe Steinhoff (see Uwe Steinhoff, “Rights, Liability, and the Moral Equality of Combatants,” *Journal of Ethics*, (2012): 339-366.

permission for the agent to resist the attack, with lethal force if necessary, if the attack infringes her right not to be killed.³⁷

On the agent-relative principle, *infringing* (as opposed to violating) the rights of someone who is not to be killed serves as a basis for an agent-relative permission to engage in defensive harm. Though killing the non-liable individual is agent-neutrally justified, the threshold at which it is permissible to kill innocents in order to promote the greater good is *lower* than the threshold at which the non-liable are required to sacrifice themselves by refraining from engaging in self-defense (which is, in turn, lower than the than the threshold at which such innocents are required to sacrifice themselves by committing active suicide).

The agent-relative principle 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. 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 liability-based justification were correct, MRTs would have no permission – not even an agent-relative one – to engage in self-defense, even assuming the agent-relative principle. But on the hybrid account, the permission to kill a lethal MRT is indeed compatible with an agent-relative permission for the threatener to engage in necessary self-defense. This is because, on the hybrid account, MRTs are not liable to be killed – rather, the permission to kill them is grounded in reasons to avert the greater evil.

The upshot is that if the agent-relative principle and the hybrid account are correct, there can be a moral “quasi-symmetry” between MRTs and their potential victims. The relationship is morally symmetrical in that both are permitted to kill the other; but the symmetry isn’t thoroughgoing since the reasons why they are permitted to kill each other are quite different, as are the constraints on the permissions. The potential victim of the MRT is permitted to kill in self-defense since killing the MRT is the lesser evil, whereas the MRT is permitted to kill in self-defense because her right not to be killed is being infringed. This stands in contrast to the liability-based

³⁷ Some might be troubled by the coherence of the claim that two parties can simultaneously bear a claim-right not to be killed and yet have a permission to attempt to kill the other. Quong defends the coherence of this view in Quong, “Killing in Self-Defense”.

justification, according to which the relationship between an MRT and her potential victim is thoroughly asymmetrical, in that only one of the two has a moral permission to engage in lethal self-defense.³⁸

In arguing for the hybrid justification for killing MRTs, I hope to have shown that the lesser-evil justification plays a far more prominent role in ‘one-on-one’ cases of defense than has been thought. Specifically, adverting to the lesser-evil justification helps makes sense of the intuition that MRTs are *victims* when they are killed – after all, they aren’t culpable for the fact-relative wrongful threats they pose. The hybrid justification, unlike the liability-based justification, accommodates this intuition by claiming that we infringe (though we do not violate) the rights of MRTs when we permissibly kill them. The result is that the hybrid justification (unlike the liability-based justification) yields the welcomed conclusions that a) multiple MRTs individually threatening one and the same innocent cannot be permissibly killed, and b) MRTs are owed compensation for the non-liable harm imposed upon them. And an implication of the claim that killing an MRT infringes her rights is that MRTs are permitted to fight back in defense against the harm permissibly imposed on them by a third party or by her potential victim.³⁹

2014

sbazargan@ucsd.edu

³⁸ Unlike others who argue for such a symmetry (see, in particular, Quong, “Liability to Defensive Harm”) the hybrid justification yields it without abandoning either the view that MRTs are morally liable to be harmed or the view that there is an agent-neutral justification for killing MRTs.

³⁹ The account has significant implications for *ex post* determinations of compensatory liability in tort law. It also has implications for the morality of war, insofar as those combatants furthering unjust aims are often only partially responsible for the threats they pose; consequently, killing them infringes their rights, given the complex account of liability. Such rights-infringements aggregate thereby making it more difficult to justify waging such a war. I leave further investigation of the practical implications of the account I’ve defended for another time.

