



Defensive Liability: A Matter of Rights Enforcement, not Distributive Justice

Susanne Burri¹

Accepted: 16 July 2021 / Published online: 30 July 2021
© The Author(s) 2021

Abstract

The Moral Responsibility Account of Liability to Defensive Harm (MRA) states that an agent becomes liable to defensive harm if, and only if, she engages in a foreseeably risk-imposing activity that subsequently threatens objectively unjustified harm. Advocates of the account contend that liability to defensive harm is best understood as an aspect of distributive justice. Individuals who are liable to some harm are not wronged if the harm is imposed on them, and liability to defensive harm thus helps ensure that harm is borne by whoever has least of an objection of justice against being burdened with it—or so advocates of the MRA insist. In this paper, I argue that liability to defensive harm is not grounded in considerations of distributive justice. While it depends on the wider societal context what allocation of a burden is distributively most just, liability to defensive harm is a distinctly localised affair. I propose that liability to defensive harm is best understood as part of the enforcement dimension of our rights to non-interference. Framing liability in this way does not invalidate the MRA. It does, however, render subjective impermissibility accounts of liability particularly promising.

Keywords Liability to defensive harm · Distributive justice · Rights enforcement

1 Introduction

May you use defensive force against a cold-blooded serial killer who is trying to strangle you? To the vast majority of people, it is obvious that you may. Moral philosophers committed to a rights-based framework provide support for this belief by arguing that the serial killer is *liable* to defensive harm. By attacking you without

For very helpful comments and discussion, I thank Lisa Hecht.

✉ Susanne Burri
susanne.burri@uni-konstanz.de

¹ Department of Philosophy, Faculty of Humanities, University of Konstanz, 78457 Konstanz, Germany

justification, the killer forfeits some of his usual rights against being harmed, and neither you nor third parties coming to your aid wrong him by using force while defending you. You also do not owe your attacker any compensation for the injuries inflicted on him. According to Jeff McMahan and other advocates of the influential *Moral Responsibility Account of Liability to Defensive Harm* (MRA), what determines a person's liability¹ is ultimately settled by considerations of distributive justice. Based on the insight that liability is closely tied to circumstances where someone's actions have made it inevitable that at least some harm will occur, they argue that we should appeal to considerations of distributive justice to settle who should be burdened with what harm, which in turn will give shape to the concept of liability.²

In this paper, I argue that this way of looking at things is mistaken. I contend that liability is properly conceived of as a matter of rights enforcement, not distributive justice. When the serial killer attacks you, you have a right to bash in his head if this is necessary to fight him off not because it is more just that he should die as opposed to you (though this may well be true). Instead, you have a right to enforce your right not to be killed, which the serial killer is threatening to violate. The enforcement dimension of your right against being killed includes an element of rights forfeiture on the part of whoever is threatening to violate it.

My main argument against the distributive justice framing of liability is that liability to defensive harm is a distinctly localised affair, whereas determinations of distributive justice are sensitive to wider societal considerations. This incongruity is removed once we regard liability as a matter of rights enforcement. Jonathan Quong agrees that we should reject a distributive justice-based framing of liability and adopt an enforcement view instead.³ In two important ways, however, Quong and I nevertheless disagree. First, unlike Quong, I do not regard it as a serious blow against the MRA that we should not assess and refine different liability accounts by appealing to considerations of distributive justice.⁴ As I see it, the MRA consists of the individually necessary and jointly sufficient conditions that it stipulates render a person liable to defensive harm. The fact that its advocates attempt to embed the MRA in considerations of distributive justice does not entail that it cannot be justified from an enforcement point of view. In fact, in one of his earliest and most detailed expositions of the MRA, McMahan leaves it open what framework is most appropriate for assessing and defending different liability claims.⁵ Second, I argue

¹ In this paper, I use "liability" as a shorthand for "liability to defensive harm."

² As McMahan puts it, "[...] the determination of liability to defensive harm is a matter of justice in the ex ante distribution of unavoidable harm." See Jeff McMahan, 'Self-Defense Against Justified Threateners', in Helen Frowe and Gerald Lang (eds), *How We Fight: Ethics in War* (Oxford and New York: Oxford University Press, 2014), 104–37, at p. 117.

³ See Jonathan Quong, *The Morality of Defensive Force* (New York: Oxford University Press, 2020), pp. 7–8; 26.

⁴ To be clear, Quong rejects the MRA not only because its advocates appeal to considerations of distributive justice to defend it. Most importantly, he argues that the MRA has intuitively implausible implications in certain cases. See Quong, *The Morality of Defensive Force*, pp. 26–9.

⁵ See Jeff McMahan, 'The Basis of Moral Liability to Defensive Killing', *Philosophical Issues* 15(1) (2005): 386–405.

that, *pace* Quong, once we consider different liability accounts from an enforcement perspective, neither the MRA nor Quong’s Moral Status Account stand out as uniquely attractive. From an enforcement perspective, the MRA and Quong’s Moral Status Account differ primarily in their conceptualisation of the circumstances under which an agent can be said to violate her victim’s rights. Which conceptualisation one deems superior depends, to some extent, on one’s intuitions about particular cases. I argue that it should also depend on the defensibility of the *control condition* the conceptualisation implies. A control condition specifies the sense in which it is necessarily under an agent’s control whether she threatens to violate another’s rights. I contend that neither the MRA nor the Moral Status Account imply a sufficiently thoroughgoing control condition. I conclude that, from an enforcement perspective and all things considered, hitherto overlooked subjective impermissibility accounts of liability deserve a closer look.

The remainder of this paper is structured as follows. In Sect. 2.1, I introduce the Moral Responsibility Account, as well as the distributive justice framing that its proponents use to support it. In Sect. 2.2, I argue that a distributive justice framing of liability should be rejected. In Sect. 3.1, I introduce and defend an enforcement framing. In Sect. 3.2, I contrast and compare leading liability accounts—including the MRA and Quong’s Moral Status Account—from an enforcement perspective. In Sect. 4, I argue that, from an enforcement perspective, we should tie an agent’s liability to circumstances that are relevantly under the agent’s control. Section 5 concludes.

2 Is Liability a Matter of Distributive Justice?

2.1 The MRA and Its Defence in Terms of Distributive Justice Considerations

According to the Moral Responsibility Account, what renders an agent liable to defensive harm is her *moral responsibility* for a threat of *objectively unjustified harm*.⁶ Harm is objectively unjustified if, in light of all the morally relevant facts—whether or not these are available to the agent at the time of acting—the balance of moral reasons speaks against the infliction of the harm. For an agent to acquire moral responsibility for a threat of objectively unjustified harm it is not necessary that they act culpably. It suffices that they voluntarily choose to engage in an activity

⁶ Key defences of the MRA can be found in McMahan, ‘The Basis of Moral Liability’; Jeff McMahan, *Killing in War* (Oxford and New York: Oxford University Press, 2009); Jeff McMahan, ‘Self-Defense Against Morally Innocent Threats’, in P. H. Robinson, S. P. Garvey and K. K. Ferzan (eds.), *Criminal Law Conversations* (New York: Oxford University Press, 2009), pp. 385–94; Michael Otsuka, ‘Killing the Innocent in Self-Defense’, *Philosophy & Public Affairs* 23 (1994): 74–94; Michael Otsuka, ‘The Moral-Responsibility Account of Liability to Defensive Killing’, in C. Coons and M. Weber (eds), *The Ethics of Self-Defense* (Oxford and New York: Oxford University Press, 2016), pp. 51–68; Kerah Gordon-Solmon, ‘What Makes a Person Liable to Defensive Harm?’, *Philosophy & Phenomenological Research* 97(3) (2018): 543–67.

where it is foreseeable that doing so might result in a threat of objectively unjustified harm. Consider the following case:⁷

The Conscientious Driver Driver is a person who always keeps her car well-maintained and always drives carefully. One day, she decides to drive to the cinema. On the way, her car veers out of control in the direction of Pedestrian, threatening Pedestrian with lethal harm. Pedestrian can avoid being run over by throwing a grenade towards the oncoming car.

The MRA implies that Pedestrian is morally permitted to throw the grenade and seriously injure or even kill Driver if this is the only way to save his life. Driver is liable to defensive harm because she knew—or should have known—that a decision to drive always carries with it a risk of objectively unjustified harm. According to the MRA, the fact that driving is associated with foreseeable harm renders it importantly different from seemingly entirely innocuous activities. While such activities might sometimes result in a threat of harm through a highly unusual chain of events, we do not render ourselves potentially liable by engaging in such activities.⁸ Consider the following case:⁹

The Cell Phone Operator Cell Phone Operator’s phone has, without his knowledge, been reprogrammed so that when he next presses the “send” button, the phone will send a signal that will detonate a bomb that will then kill an innocent person.

Cell Phone Operator is not morally responsible for the harm that he threatens because it is not foreseeable that pressing “send” on a phone might threaten serious harm to innocent others. Cell Phone Operator’s lack of responsibility entails that he is not liable to defensive harm according to the MRA.

To explain and justify the central role that foreseeable risk imposition plays within the MRA, its advocates appeal to luck egalitarian ideas of distributive justice.¹⁰ They note that, first, considerations of liability come to the fore whenever someone’s actions have made it inevitable that at least some harm will occur, and

⁷ McMahan, *Killing in War*, p. 165.

⁸ It is a difficult question whether we can draw a principled line between actions where harm is relevantly “foreseeable” and actions where no harm is foreseeable. After all, conscientious drivers reasonably believe that they impose only tiny risks of harm on innocent others. At the same time, we are all aware that seemingly innocuous activities can sometimes result in a threat of harm owing to unfortunate circumstances, which means that even innocuous activities are, in a recognisable sense, tied to tiny risks of harm as well. Elsewhere, I argue that we can meaningfully pick out foreseeably risk-imposing activities by characterising them as activities that are associated with duties of care. See Susanne Burri, ‘Morally Permissible Risk Imposition and Liability to Defensive Harm’, *Law & Philosophy* 39 (2020): 381–408, at pp. 391–6.

⁹ McMahan, *Killing in War*, p. 165.

¹⁰ See e.g. See Gordon-Solmon, ‘What Makes a Person Liable to Defensive Harm’; Otsuka, ‘The Moral-Responsibility Account of Liability to Defensive Killing’, pp. 63–4.

where we are thus faced with the question of who should bear what harm and why.¹¹ This question, they suggest, is best answered by considering what allocation of harm would be distributively most just under the circumstances.¹² Luck egalitarian theories of distributive justice claim that differences in how well people fare are not unjust if they are due to choices that we can rightly hold people accountable for. It is only when one person fares worse than another through sheer or “brute” luck that we have reasons of justice to intervene and make the worse-off person better off.¹³

Crucially for our purposes, luck egalitarians argue that responsibility for our choices can extend to the outcomes of risky choices where we did not know, when we decided to act, how things would turn out. If I invest some of my money on the stock exchange instead of leaving it in my savings account, it need not be unjust if this leaves me financially either worse off or better off than others. By investing on the stock market when I could reasonably avoid doing so, I choose to expose myself to well-defined risks. How things turn out for me is thus not a matter of what luck egalitarians call *brute luck*, but of *option luck*. Where an outcome is due to option luck, I can rightly be held responsible for it.¹⁴

Defenders of the MRA take seriously these ideas, and draw from them the implication that it is preferable, from a distributive justice perspective, that Driver rather than Pedestrian should be burdened with the harm threatened by her driving. After all, driving is a foreseeably risk-imposing activity whose dangers are well-understood, and Driver voluntarily chose to engage in it. Nothing comparable can be said of Pedestrian.¹⁵

The MRA can be criticised even if we grant that considerations of distributive justice provide the “moral foundations”¹⁶ of liability to defensive harm, and even if we concede that luck egalitarianism appropriately insists on holding people responsible for their voluntary choices. As I argue elsewhere, the idea that we should hold people responsible for their choices does not, by itself, favour allocating harm to Driver as opposed to Pedestrian.¹⁷ Holding people responsible for their choices involves clarifying, before they make choices, what potential burdens and benefits we regard as tied to particular choices, and to hold individuals to account for these burdens and benefits once the upshots of their choices become clear. This leaves open, however, what burdens and benefits we should tie to which choices and why. To settle this issue, we need to appeal to considerations beyond responsibility.¹⁸ It

¹¹ See e.g. Jeff McMahan, *The Ethics of Killing. Problems at the Margins of Life* (Oxford and New York: Oxford University Press, 2003), pp. 401–2.

¹² *Ibid.* In the passage cited, McMahan explicitly refers to the MRA as a *Justice-Based Account* (original capitalisation and italics).

¹³ See e.g. Kasper Lippert-Rasmussen, *Luck Egalitarianism* (London: Bloomsbury Publishing, 2016).

¹⁴ Ronald Dworkin, ‘What is Equality? Part 2: Equality of Resources’, *Philosophy & Public Affairs* 10 (1981): 283–345, esp. p. 293.

¹⁵ Though see Seth Lazar, ‘Responsibility, Risk, and Killing in Self-Defense’, *Ethics* 119 (2009): 699–728, at pp. 718–22.

¹⁶ McMahan, *The Ethics of Killing*, p. 402.

¹⁷ Burri, ‘Morally Permissible Risk Imposition’, pp. 403–5.

¹⁸ This point is made e.g. by Serena Olsaretti, ‘Responsibility and the Consequences of Choice’, *Proceedings of the Aristotelian Society*, New Series, 109 (2009): 165–88, and by Zofia Stemplowska, ‘Making Justice Sensitive to Responsibility’, *Political Studies* 57 (2009): 237–59.

is, moreover, far from clear that these would favour tying any driving-related harm to drivers.¹⁹ Even if they would, it would then ultimately be these further considerations, and not an appeal to responsibility, that spoke in favour of allocating harm to Driver over Pedestrian.

The just mentioned criticism is *internal* in the sense that it grants advocates of the MRA all of their assumptions. In this paper, I pursue a significantly more *external* line of criticism. I use Sect. 2.2 to argue that liability to defensive harm has its foundations not in considerations of distributive justice, but is ultimately a matter of the enforcement of rights. This point, even if it is well-taken, does not by itself speak against the MRA. After all, it is possible that the liability conditions that the MRA identifies are quite defensible from an enforcement point of view. I investigate this possibility in Sects. 3 and 4.

2.2 Against a Distributive Justice Framing of Liability

Phillip Montague seems to be the first moral philosopher who, in recent times, has framed liability to defensive harm as a matter of distributive justice²⁰ Montague argues that whenever we are faced with a “forced choice between lives,” considerations of justice dictate that we should direct harm towards those who are at fault for the forced choice situation.²¹ Roughly speaking, this is because “good things should befall those who behave well, and bad things those who behave badly.”²² Montague thus assumes a desert-based view of distributive justice, and uses this to defend a culpability account of liability.

McMahan cites Montague favourably for framing liability in terms of distributive justice, and proceeds to use the same framing to defend the MRA based on luck egalitarian underpinnings.²³ The MRA has proved extremely influential, and its advocates have unanimously accepted that an adequate defence of the account should rest on considerations of distributive justice.²⁴ Even those who reject the MRA have sometimes accepted its framing of liability in terms of distributive justice.²⁵

¹⁹ See e.g. Lazar, ‘Responsibility, Risk, and Killing in Self-Defence’, where Lazar comes close to suggesting that much more objectively unjustified harm might occur if pedestrians were to arm themselves with grenades and acted in self-defence whenever they gauged that their life was at risk due to a driving-related threat.

²⁰ Phillip Montague, ‘Self-Defense and Choosing between Lives’, *Philosophical Studies* 40 (1981): 207–19. Montague talks only about rights forfeiture and does not use the term *liability*, which is due to McMahan, ‘The Basis of Moral Liability to Defensive Killing’.

²¹ *Ibid.*, p. 215.

²² *ibid.*, p. 216.

²³ McMahan mentions Montague in *The Ethics of Killing*, p. 402, fn. 122.

²⁴ See e.g. Gordon-Solmon, ‘What Makes a Person Liable to Defensive Harm’; Otsuka, ‘The Moral-Responsibility Account of Liability to Defensive Killing’.

²⁵ See e.g. Jonathan Quong, ‘Liability to Defensive Harm’, *Philosophy & Public Affairs*, 40 (2012): 45–77. Throughout this article, Quong assumes that an account of liability is defensible to the extent that it accords with our ideas about justice. Importantly, in *The Morality of Defensive Force*, Quong no longer accepts a justice-based framing of liability.

My key claim in this paper is that framing liability in terms of distributive justice is a mistake. Liability to defensive harm is a distinctly *localised* affair; matters of distributive justice are not. There is thus a mismatch between liability and the considerations that supposedly underpin it. To see this more clearly, consider again the case of the Conscientious Driver, and assume that a luck egalitarian approach to distributive justice is essentially correct. Further suppose that attaching all driving-related costs to the choice to drive is morally most appropriate. Granting all of this, we are still not entitled to conclude that it is distributively most just if Driver, as opposed to Pedestrian, should suffer lethal harm. According to luck egalitarianism, holding people responsible for their choices is just only against a backdrop of substantive equality of opportunity. If people do not enjoy relevantly equal opportunities, equalising opportunities can trump considerations of responsibility. If, for example, Pedestrian is 85 years old and has, up to now, lived a life with ample opportunities, whereas 30 year old Driver has only just started to escape the effects of a severely disadvantaged upbringing, it may be distributively more just if Pedestrian dies.

Advocates of the MRA do not address this issue. It may be their thought that in sparsely described hypothetical cases, we always assume that “other things are held equal,” i.e. that there are no special facts about the situation not mentioned in the case that are relevant to the issue at hand. In liability cases, the idea would then be that if other things are held equal, considerations of equal opportunity do not favour any particular distribution of harm.

There are, however, two problems with this thought. First, when we assume that “other things are held equal,” we usually assume that no *special* facts obtain. In the case of the Conscientious Driver, we assume, for example, that Driver does not happen to be Pedestrian’s bodyguard who is shirking her duties to enjoy a movie at the cinema. Put differently, we grant that details left out in the description of a case are set to how things would ordinarily be. But a situation where considerations of equal opportunity do not favour any particular distribution—for example because the case is set within a society where substantive equality of opportunity is realised—is a rare exception; it is, emphatically, not the norm. If advocates of the MRA assume such a backdrop in the cases that they discuss, this needs clarification.

Second, and more importantly for our purposes, it simply seems false that the truth of our liability judgements depends on the wider societal context in the same way that judgements of distributive justice do. In my view, we can see this most clearly if we turn away from the MRA and consider instead Montague’s culpability account and the desert-based justification he provides in its favour. Suppose I am minding my own business, planting vegetables on my garden patch. A passer-by starts picking a fight with me, growing increasingly agitated by the fact that I keep ignoring him. He finally threatens to knock me into next week, and, as I continue not to respond, he storms towards me, fists raised. The fact that my attacker attacks me for no good reason suffices to render him liable to defensive harm. Even if I have done some horrible things in the past that remain unredeemed, and I thus have reason to assume that I am more deserving of what Montague refers to as “bad things” than my attacker, it seems obvious that my attacker still forfeits some of his rights against being harmed by attacking me.

I accept, of course, that my criticism of a distributive justice-based framing of liability is not decisive.²⁶ In the absence of a more suitable framing, it might thus be reasonable to continue to assess different liability accounts from a distributive justice perspective. But, as I will now argue, a more suitable framing is, in fact, available.

3 Liability as a Matter of Rights Enforcement

3.1 The Enforcement View

The key ideas of an *enforcement view of liability* are as follows. First, we all have claim-rights against being harmed, or, more generally, claim-rights not to be interfered with. With a claim-right, others owe the rightholder a duty to respect the right. It is constitutive of our rights against being harmed that they have an enforcement dimension: if they are about to be violated, there are certain things that we—or third parties acting on our behalf—are permitted to do in order to prevent the violation. Liability to defensive force is an element of this enforcement dimension. More precisely, a person who is about to violate someone else's rights thereby forfeits some of their own rights against being harmed, thus strengthening the normative standing of the person whose rights are about to be violated. In this way, the permissions we have in order to enforce our rights to non-interference go beyond the usual permissions we tend to enjoy in other circumstances.²⁷

The enforcement view of liability is not new. In recent times, Judith Thomson is the moral philosopher who has endorsed it most unambiguously.²⁸ Thomson argues that if our bodies are the cause of objectively unjustified harm to a victim, this suffices to turn us into rights violators. Threatening such harm through one's mere physical presence is thus sufficient to trigger the victim's enforcement rights and, as a conceptual flipside, our liability to defensive harm.²⁹ The enforcement view may

²⁶ For further criticism, see Quong, *The Morality of Defensive Force*, pp. 6–8.

²⁷ It is a further question whether the rights we have to aid—as distinct from the rights we have not to be interfered with—also include an enforcement dimension that entails liability to defensive harm on the part of the would-be rights violator. I answer this question in the affirmative in Susanne Burri and Lars Christie, 'On the Enforceability of Poverty-Related Duties', *Ethics & Global Politics* 12 (2019): 68–75.

²⁸ See Judith Jarvis Thomson, 'Self-Defense', *Philosophy & Public Affairs* 20 (1991): 283–310. Like Montague, Thomson does not use liability terminology in her discussion (see fn. 20). Going back considerably further in history, it is an interesting question whether John Locke's ideas about rights forfeiture are closer to an enforcement view or a distributive justice view. In his *Second Treatise*, Locke argues for a reciprocity view of rights according to which we forfeit our rights against being harmed by culpably disregarding the rights of others, thus declaring "to live by another rule than that of reason and common equity [...]." See John Locke, *Second Treatise of Government* (1690), published online at The Project Gutenberg (www.gutenberg.org/files/7370/7370-h/7370-h.htm), chs. 2 and 3, quote taken from sec. 8, ch. 2. Montague uses Locke's ideas to build his justice-based view of liability. See Montague, 'Self-Defense and Choosing between Lives', esp. pp. 213–6. Bas van der Vossen, by contrast, locates a distinctive element of rights enforcement in Locke's ideas. See Bas van der Vossen, 'Uncertain Rights Against Defense', *Social Philosophy and Policy* 32 (2016): 129–145, at p. 140.

²⁹ See Thomson, 'Self-Defense'.

be implicit in other discussions of liability as well; at times, it seems that even those who explicitly defend a justice-based view are implicitly drawn to it. Having said that, it is difficult to make such judgements with any confidence. After all, liability to defensive harm is closely tied to the rights that we have against being harmed no matter what view of its moral foundations one ultimately accepts. More precisely, liability is, by its very definition, the forfeiture of rights that we usually possess.

Thomson is a clear case of someone who accepts the enforcement view for the following reason.³⁰ Throughout her discussion of self-defence, she remains firmly focused on the potential victim who considers fighting back. Thomson's guiding question is always whether the victim retains her rights against being harmed and, in this way, her rights to enforce these rights. Crucially, the key claim of the enforcement view as I conceive of it is that *threatening a rights violation is both necessary and sufficient for liability to defensive harm*.³¹ For someone who rejects the enforcement view, threatening a rights violation and becoming liable to defensive harm may frequently coincide, but the two concepts can nevertheless come apart.³²

3.2 Leading Liability Accounts Considered From an Enforcement Perspective

Thomson's Causal View Thomson argues that if someone's physical movements threaten objectively unjustified harm, this is sufficient for threatening a rights violation and, therefore, for rendering the threatener liable to defensive harm. If, for example, I am picked up by a random gust of wind and blown towards you, I am liable to be killed if this is necessary to save your life.³³ On this view, we are thus capable of violating others' rights even when the harm that we threaten is not connected to choices that we can rightly be held morally accountable for. To many, this makes her position untenable. As Thomson anticipates, many are likely to insist that there is "[n]o violation of rights without agency."³⁴ In response, Thomson insists on her view of what it takes to violate a right based on her intuitions about particular cases. In a nutshell, it is obvious to her that if a gust of wind picks me up and threatens to hurl my body towards you with lethal force, then (i) you may kill me if this is needed to save your life; (ii) others may come to your help, killing me to save your life; (iii) I may not attempt to thwart your self-defence to save myself.

If we accept Thomson's view, our rights to non-interference provide us with extremely robust normative protections. It suffices that someone would cause us objectively unjustified harm that we are morally entitled to defend ourselves. The flip side of such strong protections is that, in our role as potential rights violators,

³⁰ Another clear case is Quong in *The Morality of Defensive Force*, where he explicitly rejects a justice-based view in favour of an enforcement view, e.g. on p. 6.

³¹ Quong seems to agree. On p. 6 of *The Morality of Defensive Force*, he argues that "[l]iability is triggered by a failure to respect other people's moral rights."

³² See fn. 42.

³³ Thomson, 'Self-Defense', p. 288.

³⁴ *Ibid.*, p. 302.

whether we respect others' rights is not something that we are necessarily in control of.

The Moral Responsibility Account Largely in response to Thomson's views, defenders of the MRA contend that an appropriate element of control is necessary for liability. They then argue that we render ourselves liable to defensive harm just in case we voluntarily perform an action where it is reasonably foreseeable, at the time of acting, that we might thereby cause objectively unjustified harm and we do, in fact, threaten such harm. As laid out in Sect. 2, defenders of the MRA justify their view by appealing to considerations of distributive justice. If we set such considerations aside and assess the MRA purely through an enforcement lens, its key appeal is that it meaningfully connects what it means to violate a right to what we do as morally responsible agents by borrowing ideas about responsibility from luck egalitarianism. Plausibly, we can respect or violate others' rights only because we are morally responsible beings who are capable of acting on reasons; a dog, for example, neither respects your rights by leaving you alone, nor does it violate them by attacking you. If so, it seems appropriate to insist that we violate others' rights only in cases where the harm that we threaten is a result of what we did as morally responsible agents.³⁵

The Moral Status Account The Moral Status Account is due to Quong, who rejects the idea that we can render ourselves potentially liable to defensive harm whenever we choose to engage in an activity that might foreseeably cause objectively unjustified harm.³⁶ To Quong, it is counterintuitive that a person such as the Conscientious Driver should be liable to defensive harm. It is even more counterintuitive that a conscientious ambulance driver who is driving to the hospital to save a patient's life—an action that the driver may well be morally required to perform based on the available evidence—should be considered similarly liable if they lose control over their vehicle.³⁷ According to the Moral Status Account, we render ourselves liable to defensive harm only by engaging in activities (i) where it is foreseeable that we might cause objectively unjustified harm by engaging in them *and* (ii) where we treat others as if they lacked certain rights that people usually possess.³⁸ Suppose, for example, that I have good reason to believe that you are culpably attacking me, when, in fact, you are not doing anything wrong. If I decide to defend myself and pull out my gun, I thereby render myself liable to defensive harm according to the Moral Status Account. After all, I treat you as if you had forfeited some of your usual rights, and I know—or should know—that my assessment might be mistaken. By contrast, conscientious drivers do not treat anyone as if they lacked rights we usually possess. Conscientious drivers' decisions to drive thus do not render them potentially liable to defensive harm.

³⁵ While it seems reasonable to defend the MRA in this way, someone drawn to Thomson's view might argue that it suffices for a rights violation that the being who threatens objectively unjustified harm *is* a moral agent, even if the threat that they pose is not meaningfully connected to their agency.

³⁶ Quong, 'Liability to Defensive Harm' as well as Quong, *The Morality of Defensive Force*.

³⁷ Quong, *The Morality of Defensive Force*, pp. 26–9.

³⁸ I gloss over some technicalities here. See Quong, *The Morality of Defensive Force*, pp. 34–45. Quong's precise definition of liability can be found on p. 38.

From an enforcement perspective, the Moral Status Account sets the bar for what it means to violate a person's right even higher than the MRA does. It thus further weakens the extent to which our rights protect us as rightholders, while strengthening the extent to which it is under our control whether we violate others' rights. In Quong's view, his account implies that agents can become liable to defensive harm only in cases where "we can reasonably demand that they refrain from performing the act."³⁹ This, it seems to me, is incorrect. If all the available evidence (falsely) suggests that a serial killer is about to attack me, it does not seem reasonable to demand that I refrain from thwarting the perceived attack.⁴⁰

Culpability Accounts Different culpability accounts all share the assumption that we become liable to defensive harm only by acting culpably.⁴¹ From an enforcement perspective, this implies that we violate others' rights only through our culpable actions. Culpability accounts thus put it firmly within our control whether we violate others' rights: violating a right is always something that we ought to avoid and that we are blameworthy for if we go through with it. In this way, culpability accounts set the bar for rights violations as high as is reasonably defensible, and equip us with comparably weak normative protections in our role as rightholders.

Culpability accounts are not entirely compelling from an enforcement perspective because common language use suggests that we can violate others' rights also by acting non-culpably. If someone holds a gun to my head and orders me to kill an innocent person, it seems that we would want to say that I violate my victim's rights, even if I am excused and thus lack culpability for doing so. Of course, those attracted to a culpability account of liability are free to insist that the only person who is a rights violator is the person threatening me, and that even though I execute the killing, I am merely the means, so to speak, through which the rights violator acts.⁴²

My presentation of leading liability accounts in this section leaves it unclear which account we should endorse from an enforcement perspective. In the next section, I will argue that considerations of appropriate control push us a fair way towards the culpability end of the liability spectrum.

³⁹ Quong, *The Morality of Defensive Force*, p. 6.

⁴⁰ A demand to refrain might of course be reasonable if the misunderstanding can be cleared up. Quong's view implies, however, that a demand to refrain would be reasonable even if the person issuing the demand knew no more than I did.

⁴¹ Culpability accounts are defended e.g. by Montague, 'Self-Defense and Choosing between Lives'; Kimberly Kessler Ferzan, 'Justifying Self-Defense', *Law & Philosophy* 24 (2005): 711–49, and 'Culpable Aggression: The Basis for Moral Liability to Defensive Killing', *Ohio State Journal of Criminal Law* 9 (2012): 669–97. Before endorsing the Moral Responsibility Account, McMahan defended a culpability account as well. See Jeff McMahan, 'Self-Defense and the Problem of the Innocent Attacker', *Ethics* 104 (1994): 252–90.

⁴² Another possibility is to reject the enforcement view that I defend in this paper, and with it the close connection between threatening a rights violation and rendering oneself liable to defensive harm. This is what Ferzan does by defending a culpability account of liability, while at the same time assuming that we can violate others' rights non-culpably. See Ferzan, 'Justifying Self-Defense', p. 714.

4 Why Rights Violations Should be Subject to a Thoroughgoing Control Condition

As the discussion in Sect. 3 has made clear, different liability accounts trade off in different manners the normative protections that we enjoy as rightholders, and the control that we have over our moral status in our role as would-be right violators. At one end of the spectrum, Thomson's causal account maximises the former, giving no weight to the latter. At the other end of the spectrum, culpability accounts limit the protections we enjoy as rightholders, but provide us with significant control in our role as would-be right violators.

In this section, I will argue that providing us with significant control in our role as would-be right violators is an important consideration that any plausible account of liability has to accommodate. My argument in favour of this claim runs as follows. If someone threatens to violate a right, this fact lowers their moral status so that the right that is threatened can be enforced. To have one's moral status lowered is normatively highly significant, and it is thus something that we should be able to justify to the person whose status is lowered. More precisely, we should be able to offer an explanation why it is morally more appropriate that they should be harmed as opposed to the person whom they threaten, as the enforcement mechanism of our rights to non-interference implies. It is not very satisfactory if we can offer no more, in terms of a justification, than "Your moral status is lowered because you are threatening to violate a right, though we acknowledge that the reason why you are in this unfortunate situation is due entirely to your bad brute luck." Intuitively, it is objectionable if someone's moral status can be lowered—such that they are no longer owed the consideration that we are usually due—based on circumstances that are beyond their control. On Thomson's causal view, however, such an objectionable state of affairs is precisely what we will be stuck with at times. Admittedly, if I do not become liable to defensive harm if a gust of wind picks me up and hurls me towards a potential victim, and if, as a result, my victim dies, then my victim's death is due entirely to bad brute luck as well.⁴³ Crucially, however, the victim dies without first having her moral status lowered; it is thus *not* morally more appropriate that my victim should die and that I should live. If my victim dies because I am non-liable, her death is not due to a normative asymmetry between the two of us, but only to differences in our non-normative circumstances.

On culpability views, we enjoy relevant control over our moral status in our role as would-be right violators. If I culpably threaten objectively unjustified harm, others can explain to me that my moral status is lowered because I am acting in a way that is morally prohibited, and that, for this reason, a morally protected good is now at risk. My moral status is lowered to protect what would not be at risk had I acted

⁴³ Even if I am not liable to defensive harm, my victim may still be morally permitted to defend herself based on an agent-relative prerogative. Third parties would then not, however, be permitted to intervene on the victim's behalf, and if she can be saved only through third-party intervention, she will die due to bad brute luck if I happen to crush her. Quong defends agent-relative prerogatives to save one's own life in situations such as the one under consideration here in Jonathan Quong, 'Killing in Self-Defense', *Ethics* 119 (2009): 507–37.

in accordance with what was morally required of me. I thus have no valid complaint against having my status lowered. This is a very compelling justification, and one that helps explain why most accounts of liability agree that even if culpability is not necessary for liability, the extent to which a person is liable to defensive harm increases with her culpability.⁴⁴

It is possible, however, that the just proffered explanation is stronger than what is needed to justify to a would-be rights violator why their moral status is lowered. After all, both the MRA and the Moral Status Account introduce an element of control in their implied definition of what it means to violate a right. For one thing, both accounts imply that it is always epistemically accessible to me, at the time of acting, that an action that I choose to perform might violate others' rights. Moreover, the MRA borrows an influential notion of control from luck egalitarianism, and one might hope that if this notion fits the bill in the context of distributive justice, it does so also in a rights enforcement context. The Moral Status Account strengthens the notion of control that the MRA introduces without, however, going so far as to make culpability a necessary condition of liability.

In my view, the MRA and the Moral Status Account take seriously, *but not sufficiently seriously*, the idea that our moral status should not be lowered due to circumstances beyond our control. The control that the two accounts introduce is not sufficiently thoroughgoing because there is no clear sense in which an agent ought to avoid all actions that might render her liable to defensive harm. To see what I mean, note that both accounts imply that we can become liable to defensive harm in cases where we acted morally permissibly—or possibly even in ways that were morally required—in light of the evidence that was available to us at the time of acting. But if I acted in a way that was entirely defensible in light of what I knew and should have known, it is difficult to see how it could nevertheless be appropriate to ground any liability to defensive harm in my actions. If this is right, then we should never be taken to violate others' rights if we acted subjectively entirely permissibly, i.e. permissibly both in light of our beliefs and our evidence.

This does not mean, however, that we ought to accept a culpability account. As I have briefly laid out in Sect. 3.2, common language use suggests that we can violate others' rights both culpably and non-culpably. If we factor this in, this suggests that a *subjective impermissibility account of liability* might be most appropriate. A suitable such account might stipulate that we become liable to defensive harm just in case we act impermissibly in light of the available evidence or in light of our beliefs, and our actions subsequently threaten objectively unjustified harm.⁴⁵ Consider again the situation where someone holds a gun to my head and orders me to kill an innocent

⁴⁴ See, for example, Jeff McMahan, 'The Ethics of Killing in War', *Ethics* 114 (2004): 693–733, at p. 724.

⁴⁵ Subjective impermissibility accounts—or evidence-relative accounts, as they are usually called—are underdefended in the literature; I am currently working on a paper defending one such account. Evidence-relative accounts are mentioned favourably in Bradley J. Strawser, 'Walking the Tightrope of Just War', *Analysis* 71 (2011): 533–44. For an evidence-relative account that *exacerbates* the control issues discussed in this paper, and that I deem implausible for this reason, see van der Vossen, 'Uncertain Rights against Defense'.

person. If I obey the order, I act under duress, and what I do may be fully excused and, therefore, non-culpable. What I do is nevertheless impermissible in light of the available evidence, and if I am morally decent, I will be painfully aware of this fact. That my actions should trigger the enforcement rights of my would-be victim is thus not something that should strike me as objectionable. Of course, it is deeply troubling that I am at the mercy of a maniac. But that my attack on an innocent person—an action that is recognisably seriously wrong—may rightly be thwarted is not similarly disturbing. In fact, it is plausibly part of the wrong committed by the maniac that I am coerced into harming an innocent person, thus rendering myself liable to defensive harm and, should my action succeed despite this fact, condemning me to live with the knowledge that I have wrongly taken another's life.

While considerations of control push us towards accepting a subjective impermissibility account along the lines just sketched, such an account has counterintuitive implications in some cases. Consider the following:⁴⁶

Duped Soldiers A group of young soldiers is successfully fooled by a totalitarian regime into believing that the regime is good and just and is under repeated attacks from their evil neighbours, the Gloops. The regime's misinformation campaign is subtle and convincing—the soldiers are justified in believing what they are told by the regime. Once the misinformation campaign is complete, these soldiers are given orders to attack and destroy a Gloop village on the border which, they are told, is really a Gloop terrorist camp plotting a major attack. In fact, everything the regime has said is a lie, and the Gloop village contains only innocent civilians. The soldiers prepare to shell the village and are about to (unknowingly) kill all the innocent civilians in it. A peacekeeping force from a neutral third country patrols the border and could avert the attack, but only by killing the soldiers.

According to the subjective impermissibility account just sketched, the soldiers are not liable to defensive harm, and the peacekeeping force is morally required to refrain from interfering. I agree with Quong that this is counterintuitive. At the same time, I fear that such counterintuitive results are the price we have to pay if we want to justify to those whom we hold liable why the moral rules favour harm to them over harm to their victims.

In the end, I doubt that a wholly satisfactory account of liability—i.e. an account that is simple, well-supported by arguments from an enforcement perspective, and, moreover, in accordance with our intuitions about a range of tricky cases—is within our reach. While I find it counterintuitive that the peacekeeping force may not proceed in *Duped Soldiers*, my intuitions are insufficiently strong to push me towards accepting a weaker notion of control such as the one implicit in the Moral Status Account. Quong doubtlessly disagrees.

⁴⁶ Quong, *The Morality of Defensive Force*, pp. 33–4.

5 Conclusion

In this paper, I have argued against the distributive justice-based framing of liability that advocates of the MRA use to defend their view of liability. I have argued that we should reject a such a framing because liability is a distinctly localised affair, whereas distributive justice is sensitive to the wider societal context.

The localised nature of our liability judgements makes sense once we switch from a distributive justice-based framing to an enforcement view. This view proceeds from the insight that our rights to non-interference all come with an enforcement dimension: if they are about to be violated, there are certain things that we—or third parties acting on our behalf—may do to prevent their violation. Defensive liability is then conceptualised as the flip side of the enforcement rights that a rightholder has against her would-be rights violator.

From an enforcement perspective, no liability account emerges as uniquely preferable. But if we take seriously the idea that we can become liable to defensive harm only in circumstances where such a reduction in our moral status is relevantly under our control, subjective impermissibility accounts of liability deserve a closer look.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.