

WHAT MAKES A PERSON LIABLE TO DEFENSIVE HARM?

Kerah Gordon-Solmon

Queen's University

ABSTRACT

On Jeff McMahan's influential 'responsibility account' of moral liability to defensive killing, one can forfeit one's right not be killed by engaging in an ordinary, morally permissible risk-imposing activity, such as driving a car. If, through no fault of hers, a driver's car veers out of control and toward a pedestrian, the account deems it no violation of the driver's right to save the pedestrian's life at the expense of the driver's life.

Many critics reject the responsibility account on the grounds that, first, it has draconian implications for threateners like the driver, and second, it contravenes the plausible principle that wronging one's victim is necessary for forfeiting one's rights. But I argue, drawing on the account's luck-egalitarian underpinnings, that (1) the account lacks the draconian implications widely attributed to it, and (2) contrary to what many assume, wrongdoing is unnecessary for rights-forfeiture. Via these arguments, I seek both to deepen our understanding of the responsibility account, and to reissue it in a more plausible and attractive form.

INTRODUCTION

Conscientious Driver: A person who keeps her car well maintained and always drives cautiously and alertly decides to drive to the movies. Freak circumstances cause the car to go out of control. It has veered in the direction of a pedestrian whom it will kill unless she, or a third party, blows it up with a grenade.¹

Driving carefully under normal circumstances is a morally permissible type of activity. The driver is not blameworthy, nor at fault. Nevertheless, driving, even when it is pursued exemplarily, is a type of activity that foreseeably exposes oneself to, and imposes on others, a slight risk of severe harm. Since the driver chose to impose that risk, which will now eventuate in someone being killed, and since the pedestrian, by stipulation, made no comparable choice, the driver is morally liable to be killed to save the pedestrian.²

This is the flagship argument with which Jeff McMahan advances his ‘responsibility account’ of moral liability to defensive harm (henceforth “liability”). The account identifies certain facts about choice and responsibility, on which judgments of distributive fairness are predicated, as also forming the basis of liability. The driver’s choice to engage in a risk-imposing activity makes her potentially liable to some degree of defensive harm. Since the driver’s choice threatens to kill another, innocent person without justification, and since that threat will be averted just in case the driver is killed, the driver is in fact liable to be defensively killed. By the standard definition of liability,

¹ Jeff McMahan, “The Basis of Moral Liability to Defensive Killing,” *Philosophical Issues*, Vol. 15 (2005), p. 393 and *Killing in War*, (Oxford: Oxford University Press, 2009), p. 165. (I’ve blended the two iterations of the example).

² McMahan (2009), p. 177.

the driver has *forfeited her right* not to be killed in the pedestrian's defense; she is not wronged by, and has no right to defend herself against it.

In his first major presentation of the account, McMahan takes the judgment that the driver is liable to be killed as datum; he defends the responsibility account as providing a compelling argument – to wit, the one summarized above – to undergird it.³ That the account supports this judgment is uncontroversial. But many critics regard this support, not, as McMahan does, as a virtue of the account,⁴ but, on the contrary, as a powerful strike against it. These critics advocate alternative accounts of liability, which they seek to defend by showing them to avoid the responsibility account's draconian implications.

My principal aim in this paper is to defend the responsibility account against this prevalent line of attack. To this end, I employ, as the account's juxtapositional foil, Jonathan Quong's 'moral status account'.⁵ I focus on Quong's alternative, not only because the arguments for it appear persuasive, but also, and more importantly, because scrutinizing those arguments affords new insight into the base commitments of the responsibility account, and into what is (and is not) entailed by those commitments. I believe, and over the main arguments of the paper shall demonstrate, first, that the responsibility account has unexploited potential to supply deeper moral arguments for its attributions of liability across a range of cases, and second, that these arguments entail greater limits on the driver's liability than the account is normally thought to support.

³ McMahan (2005), pp. 393-396.

⁴ McMahan (2005), p. 398.

⁵ Jonathan Quong, "Liability to Defensive Harm" *Philosophy and Public Affairs*, Vol. 40 No. 1 (2012), pp. 45-77.

I shall defend two main conclusions about the driver case, the first of which is party line: it says the responsibility account compellingly supports the judgment that the driver is liable to lethal defensive harm. The second conclusion says, *pace McMahan et al.*, that the account does not support, but in fact contravenes the judgment that the driver is liable to having that harm inflicted on her by means of throwing a grenade at her. In defending the latter, I shall draw upon the account's employment, first, of luck-egalitarian reasoning, which I will show to have subtler implications in the domain of preventive justice than has been recognized, and second, of certain basic nonconsequentialist precepts, the implications of which, in this domain, have not yet been unpacked.

The significance of the paper's arguments belies the relative modesty of its express conclusions. What will emerge from this study is a radical reinterpretation of the responsibility account. A superior rationale for it will be advanced, and its load-bearing parts will be restructured in the form of a systematic theory. This study will also challenge the rare orthodoxy in recent self-defense literature, namely, that wrongdoing is necessary for liability. The luck-egalitarian underpinnings of liability provide decisive reason to reject this postulate. Or so I shall argue.

The paper is organized as follows. Section I reconstructs the responsibility account in its canonical form; it foregrounds those features of the account that I will deploy in subsequent sections. Section II dismantles Quong's arguments for his moral status account; and it defends, against Quong, the basis for the judgment that the driver is liable to lethal defensive harm. Section III responds to the criticism that the responsibility account is draconian; it argues the account is resourced to oppose the claim

that the driver is liable to be defensively killed via the pedestrian's grenade. Section IV reappraises the canonical responsibility account; it identifies a disparity between the canonical account and the version I defend, and it argues that the latter is superior.

I

On the responsibility account, it is a necessary condition for liability to defensive harm that a person be *morally responsible* for a threat of *unjustified harm* to another.⁶ The task of this section is to unpack this claim, and, in so doing, to explicate certain features of the responsibility account that will be important to the argument of the present paper. I'll first briefly address the matter of justification, then turn to the matter of moral responsibility as it features in the account. In addressing the latter, I'll speak, first, to moral responsibility as the basis of liability; second, to the relevant factors that increase or reduce the degree of person's responsibility and thereby the extent of her liability; and third, to the two distinct ways in which the extent of a person's liability can be increased or reduced. I'll appeal to these three points in section III, in arguing that the responsibility account supports greater restrictions on the driver's liability than it is normally thought to support. I'll return to the matter of justification in section IV; there I'll argue, appealing to theoretical points I'll make in section III, that justification does more load-bearing work in the responsibility account than most people recognize; specifically, I'll argue that it renders other, seemingly important, but, I'll claim, problematic, features of the account superfluous.

⁶ McMahan (2009), pp. 34-35, and p. 157.

i. Justification

A number of candidate justifications for harming another have been defended within the contemporary literature; here's an indicative list. An inflicted harm can be justified by the victim's being liable to that harm;⁷ or consenting to, deserving, or having an enforceable duty to bear that harm; or it can be justified by the threatener's having an agent-relative prerogative, or an associative duty (to some third party, such as a family member or compatriot) to inflict that harm. It can also be justified as the lesser evil, if the victim's being harmed in the relevant way is significantly less impersonally bad than the available alternatives.

What's important for the present purposes is that, by harming another without justification, a person incurs liability to defensive harm, and, conversely, by harming another with justification (of which there are several types), a person insulates herself against liability to defensive harm.⁸

ii. Moral Responsibility

On the responsibility account, a person is morally responsible for a threat of harm just in case she chooses to engage in a foreseeably risk-imposing activity, the risk of which will eventuate in harms.⁹ Driving is a paradigm example of a foreseeably risk-imposing activity, and so the choice to drive renders the conscientious driver morally responsible for the harms that eventuate from her driving. Likewise, though it is unnatural to put it this way, attacking someone with the intent to do bodily harm is a

⁷ McMahan calls threateners who act with liability justifications not only justified, but *just* threateners. Jeff McMahan "Self-Defense Against Justified Threateners," in Helen Frowe and Gerald Lang, eds., *How We Fight: Issues in Jus in Bello* (Oxford: Oxford University Press, 2014), p. 107.

⁸ For defense of the latter claim, see McMahan (2014).

⁹ McMahan (2005), p. 394.

foreseeably risk-imposing activity. In either case, if the threatened harm is unjustified, and defensively harming the threatener will be effective in averting the threat, and is necessary for doing so, the threatener is presumptively liable to defensive harm.

As indicated earlier, the responsibility account is undergirded by certain intuitions about distributive fairness:¹⁰ other things equal, if I voluntarily engage in a risk-imposing activity that will eventuate in harms, then I ought to be the one to suffer those harms.

In its appeal to distributive fairness, the responsibility account of liability is the parallel, in the domain of preventive justice, to luck egalitarianism in the domain of distributive justice. I'll have more to say about this parallel (and in particular, about its implications with respect to the driver case) later. For now, let me make its structure explicit.

Luck egalitarians claim that it is unfair, and therefore unjust, for one person to be worse off than another through no fault or choice of her own. If a person's misfortune is the result of bad brute luck (if, for example, she's the victim of a falling meteor), she has a fairness-based claim against others for compensation. By contrast, if a person's misfortune is the result of bad option luck (if, for example, she invested in stocks, the value of which subsequently fell), she has no such claim against others for compensation. There is no unfairness in requiring people to shoulder the costs of their bad option luck.

In ordinary distributive-justice cases, when a person takes a gamble, she forfeits certain claims against others for assistance if that gamble turns out badly. In parallel, in self-defense cases, when a person engages in a risk-imposing activity and the risk is realized, she forfeits certain claims against others not to harm her in the course of

¹⁰ Quong (2012), p. 56.

avoiding the harms for which she is responsible.¹¹ Just as it would be unfair to redistribute the costs of one's financial gamble onto others, so it would be unfair to allow the harms resulting from one person's risk-imposing activity to be inflicted on others. They ought instead to go to the risk-imposer, via defensive action.

iii. Other Factors Influencing Moral Responsibility, and Hence, Liability

The choice to engage in a risk-imposing activity, the risk of which will eventuate in an unjustified harm to another, is, on the responsibility account, a threshold condition for moral responsibility for that harm, and thus for liability to defensive harm. Above that threshold, the degree of a person's moral responsibility tracks various additional factors, and the extent of the harm to which she is liable varies, other things being equal, with the degree of her moral responsibility.

Of the factors that affect liability, I shall discuss four.

First, the extent of a person's liability will be affected by the magnitude of the harm she threatens: other things equal, the greater the magnitude of the threatened harm, the greater the potential liability of the threatener. An unjustified threatener is potentially liable to up to at least the magnitude of harm with which she threatens her victim.

The second factor affecting the extent of a person's potential liability is whether, and to what extent she is blameworthy, or culpable for her threat of unjustified harm.

¹¹ A victim assuming certain risks can, correlatively, insulate her threatener against liability to being defensively harmed on her behalf. Consider a NASCAR spectator who intentionally crosses a safety barrier to get a better view of the track. No racing driver is liable to be defensively killed to save the barrier-crosser from her out-of-control vehicle. This is because, in the context of the race (relevant features of which include the publicized function of the barriers to keep everyone – spectators and drivers – safer), the barrier-crosser's voluntary act increases the risk of *someone* present being seriously harmed. In such cases, the risk-assumer is best characterized as a special type of risk-imposer. I'm grateful to Saba Bazargan for discussion of this point. The example belongs to Victor Tadros, *The Ends of Harm* (Oxford: Oxford University Press, 2011), p. 231.

The extent of a person's liability increases with her degree of blameworthiness. For example, a wholly blameworthy attacker threatening her victim with a broken arm would potentially be liable to the more serious harm of a broken leg: her culpability makes the greater harm proportionate. A person is blameworthy if she threatens harm recklessly, negligently, or intentionally, without excuse.¹²

The third factor influencing a person's liability is whether, and to what extent, she is excused. Excuses mitigate a person's blameworthiness, and, correspondingly, decrease the degree of harm to which she is potentially liable. If a person is fully excused, then she is completely blameless. If she is only partially excused, then she is blameworthy to the extent to which she lacks excuse.¹³ A person may be excused if, for example, she is operating on the basis of false information, or under duress.

Nonconsequentialist considerations constitute a final factor affecting the extent of a person's liability. It makes a difference to someone's liability whether the unjustified harm she threatens is intended or merely foreseen, whether it will be inflicted as a means or as a side effect, and so forth. It is more seriously wrong, for example, to harm a person intentionally than it is to harm a person foreseeably but unintentionally; and it can be more seriously wrong to harm a person foreseeably than to take a known risk of harming a person and, through bad luck, harming her accidentally.¹⁴ A person thus has a higher degree of liability for her intentional than for her merely foreseen threats of harm, and can likewise have a higher degree of liability for her foreseen than for her accidental threats of harm.¹⁵ Other things equal, a person's right against being intentionally harmed

¹² McMahan (2009), p. 159.

¹³ McMahan (2009), pp. 159-162.

¹⁴ But not always; exceptions include accidental harms caused by negligence.

¹⁵ McMahan (2009), p. 167.

is more difficult to defeat, or is more stringent, than her right against being foreseeably but unintentionally harmed, and the more difficult a right is to defeat, the greater one's potential liability when one contravenes it without justification. This claim does not speak to all accidental harms: it strikes me as implausible that a person can have a right against being killed accidentally, *simpliciter*.¹⁶ But this just makes explicit what has so far been tacit, namely, that contravening a person's right (or wronging a person) is not a necessary condition for liability on the responsibility account.¹⁷ I will defend this feature of the account in section II, and cash out its implications in section III.

iv. Mitigated Liability

So far, I have treated the question of the extent of a person's liability as the question of the ceiling on a person's liability – as the question of the maximal harm to which a person is potentially liable, given the nature of her threat. But there's a second sense in which the extent of a person's liability can be affected. To see this, contrast the conscientious driver with a paradigmatic culpable threatener: a man who threatens to murder his wife so he can inherit her money. Both the driver and the husband are morally responsible for lethal threats, and so both the driver and the husband are potentially liable to lethal harm. But while the driver is blameless for threatening to kill a person accidentally, the husband is fully to blame for threatening to kill a person

¹⁶ My underlying supposition here (which I share with Quong) is that a person has rights against others that they *treat her*, or *refrain from treating her* in certain ways, but does not have rights against *having things happen* to her. Being accidentally harmed as a result of another's action is insufficient evidence that one has been mistreated, and so insufficient evidence that one's rights have been violated or infringed.

¹⁷ McMahan would reject this characterization of the account. I'll discuss his reasons for doing so, and why I find them unpersuasive, in section IV.

intentionally. So the degree of the husband's responsibility, and thus the extent of his liability, will be greater than the driver's.

The apparent tension between these claims is resolved as follows. Both the driver and the husband are liable to fatal harm, if that harm is necessary to avert their respective threats. But suppose the pedestrian and the wife each can defend her life effectively in either of two ways: each can kill her threatener, or can inflict on her threatener serious but non-fatal injuries, at the cost of suffering substantial injury herself (i.e., a broken leg). In these circumstances, the driver is only liable to the latter harm: an implication of the lesser extent of the driver's liability is that the pedestrian must allow herself to suffer non-negligible, even substantial harm, in order to avoid killing the driver, assuming that option is available.¹⁸ The husband, however, remains liable to be killed: the relevant factors exacerbate rather than mitigate the extent of his liability; and so the wife is not obliged to allow herself to suffer substantial harm in order to spare him.

I'll discuss how this understanding of mitigated liability reconciles with the luck-egalitarian underpinnings of the responsibility account in section III.

II

Rejecting the responsibility account is normally concomitant with rejecting luck-egalitarian distributive principles as the basis of liability.¹⁹ In this section, I shall defend

¹⁸ McMahan calls this the "morally best" available option; he argues the necessity condition of liability is only satisfied if this option is chosen. Jeff McMahan, "The Limits of Self-Defense," in Christian Coons and Michael Weber, eds., *The Ethics of Self-Defense* (New York: Oxford University Press, 2016), pp. 187-188.

¹⁹ Saba Bazarban excepted; see "Killing Minimally Responsible Threats," *Ethics* Vol. 125 No. 1 (2014).

the relevance of distributive principles to preventive justice, via critical scrutiny of Quong's moral status account. My argument here will be negative: it will show that, contrary to appearances, anti-distributive theories of liability do not claim the *de facto* allegiance of our theoretical-level intuitions; nor have superior arguments been made on their behalf. Their seeming advantage over the responsibility account derives only from their palatable conclusions about the driver case. Hence, in section III's positive defense, I shall reissue the responsibility account's treatment of the case, softening its implications.

On Quong's moral status account, a person is liable to defensive harm just in case she fails to treat others in accordance with the moral claims they in fact exert on her, and, in so doing, threatens an innocent (that is, nonliable) victim. A person is not liable to defensive harm if she treats others in accordance with the moral claims they in fact exert on her, or if, through good luck, her wrongful conduct does not threaten others.

Quong illustrates his view with two paradigm cases in which a person makes herself liable to defensive harm. The first is one of reckless driving. Someone decides to drive to the movies, but finds herself running late, and, fearing she'll miss the beginning of her film, drives recklessly, at breakneck speed. She loses control of the car and careers in the direction of a pedestrian. The reckless driver is acting impermissibly: she is exposing others to an unduly high level of risk in order to secure a trivial benefit for herself. On this occasion, her wrongful conduct will result in an innocent person being killed unless she is killed in that person's defense. The reckless driver is therefore liable to be killed: by treating others in a way that contravenes their moral claims against her, when doing so will eventuate in their harm, she forfeits certain claims she would

otherwise have against them not to be harmed in their defense (to at least the degree to which she would otherwise harm them).²⁰

Michael Otsuka's "dignitary" is a second paradigm case:

Imagine that you extend your hand to shake the hand of some foreign dignitary at a reception. Unbeknown to you, a third party projects a stunningly realistic holographic image of a pistol onto your hand. The dignitary, who is accustomed to threats on her life, sees the hologram, forms the justified belief that you are about to assassinate her, and coolly draws a pistol in order to shoot you down in self-defense.²¹

Unlike the reckless driver's, the dignitary's conduct is not blameworthy. The best available evidence overwhelmingly supports the dignitary's belief that you are culpably threatening her life, and thus, that you have forfeited your moral claim against her not to be lethally shot. But the dignitary's belief that you've forfeited this moral claim against her is false. In acting on this belief, the dignitary treats you as *lacking a moral claim against her that you in fact possess*. She therefore makes herself liable to being defensively killed.

Quong contrasts these two types of cases with that of the conscientious driver. The reckless driver and the dignitary both treat others as lacking moral claims they in fact possess. The reckless driver does so through her general mode of conduct. The dignitary does so when she acts on the basis of a fatal, if faultless, mistake that she makes about the "moral status" of her victim. The conscientious driver's conduct, however, is not

²⁰ Quong (2012), p. 65.

²¹ Michael Otsuka, "Killing the Innocent in Self-Defense," *Philosophy & Public Affairs*, Vol. 23 No. 1 (1994), p. 91.

predicated on anyone lacking moral claims against her they in fact possess. Quong thus contends that the conscientious driver is immune from liability.

The distinction Quong draws between the reckless driver, the dignitary, and the conscientious driver can be restated as follows (although Quong does not put it this way himself). The reckless driver impermissibly risks accidentally harming an innocent person. The dignitary permissibly risks intentionally harming an innocent person. The conscientious driver permissibly risks accidentally harming an innocent person.²² Both the reckless driver and the dignitary wrong their victims; the reckless driver by imposing on her victim an impermissible risk of harm; and the dignitary by intentionally killing her victim when she is in fact innocent. The conscientious driver, by contrast, does not wrong anyone with her conduct. Since, for Quong, both wronging and threatening harm are necessary conditions for liability to defensive harm, the conscientious driver is not liable.

Formulating Quong's view in this way telegraphs his response to the following objection (which is motivated from within the responsibility account). On the face of it, the conscientious driver and the dignitary are symmetrical in the following respect. The driver and the dignitary each *runs the risk* that she will fatally harm an innocent person. We can stipulate that the odds in each case are the same – that is, that the odds that someone in the driver's circumstances will kill a bystander are the same as those that someone in the dignitary's situation will kill a nonliable victim. Given that the driver and the dignitary each runs the same risk of fatally harming an innocent person and that both their risks will now eventuate, it seems unfair that only the dignitary and not the driver should be liable to defensive harm.

²² I'm using "permissible" in these sentences in the evidence-relative sense.

Quong's response to this objection is that only the dignitary risks *wronging* the nonliable person she harms – a difference that suggests an asymmetry in liability between the dignitary and the driver. But this response does not fully dispense with the objection, which Otsuka has reissued, revised as follows: Quong has identified an asymmetry between the dignitary and the driver, which supports their asymmetrical liability. But Quong has not shown the driver to be immune from liability, since there remains a plausible alternative; namely, that the driver is liable to a lesser extent (or to a lesser degree of harm) than the dignitary.²³ So Quong has not shown that wronging one's victim is a necessary condition for liability, rather than one factor in determining the amount of harm to which a threatener is liable, as the responsibility account dictates.

To defend his claim that the driver is not liable, Quong says the following:

If we treat others as if they are liable to harm, it seems only fair to suppose that we may become liable to defensive harm should that judgment be mistaken. But when our risk-imposing actions do not treat anyone else as lacking moral rights against harm..., it would be unfair and wrong to hold us liable to defensive harm if our actions unfortunately threaten harm against innocent others.²⁴

Quong is not appealing here to fairness in its distributive (or telic) sense, as proponents of the responsibility account do. Instead, he is appealing to the idea of fair treatment (or to fairness in its deontic sense). The intuitive claim Quong makes with respect to the driver is this: that it's unfair to treat a person any worse than she treats others. Since the driver does not treat others as lacking moral claims they in fact possess, it would be unfair to

²³ Michael Otsuka, "The Moral Responsibility Account of Liability to Defensive Killing," in Christian Coons and Michael Weber, eds., *The Ethics of Self-Defense* (New York: Oxford University Press, 2016), pp. 66-67.

²⁴ Quong (2012), p. 69.

treat *her* as lacking moral claims.

The intuitive claim that it is unfair to treat a person worse than she treats others strikes me as uncontroversial. But, contrary to appearances, it does not support the inference that it would be unfair to treat the driver as lacking any moral claims. The intuitive claim, unpacked, says that if a person (i.e., the driver) does not treat others as lacking moral claims they in fact possess against her, then it would be unfair to treat her as lacking moral claims she in fact possesses against them. It thus leaves open the question of *which* moral claims the driver in fact possesses against others. It leaves open the question of whether the driver has in fact forfeited any of her moral claims.

If Quong intends to argue for his account of liability from uncontroversial intuitions about fair treatment, his argument is circular: it assumes what it purports to prove, namely, that treating others as lacking moral claims they in fact possess is a necessary condition for forfeiting any of one's own moral claims. I suspect what Quong instead has in mind is a robust conception of fair treatment, according to which it is unfair to treat a person as lacking *any* moral claims just in case she does not treat others as lacking moral claims they in fact possess. But this is just to restate the basic tenet of Quong's moral status account in fairness terms; it cannot serve in its defense.

Quong's proposal is attractive, but, in the absence of a positive argument in its favour, which Quong does not provide, I think it should be rejected in favour of the following Otsuka-like alternative: that *if a person treats others in accordance with the moral claims they in fact exert against her, this limits the extent to which she can forfeit her moral claims against others* (and thus, limits the extent of her potential liability to defensive harm).

One reason I think Quong's principle should be rejected is that I think the luck-egalitarian principle is true – and Quong's proposal is incompatible with the luck-egalitarian principle. That principle, remember, says people have fairness-based claims to compensation for the costs of their bad brute luck, but not for the costs of their bad option luck (i.e., of their optional risk-taking). It thus implies the following:

If a person takes an optional risk that eventuates in a loss, she (thereby) forfeits her fairness-based claim to be compensated for that loss – a claim which she would retain if, counterfactually, the loss were the product of bad brute luck.

I take it as given that (1) a person can take an optional risk that turns out badly without thereby treating others in disaccord with the moral claims they in fact exert against her, and (2) fairness-based claims are a type of moral claim. Thus, according to the luck-egalitarian principle – and *contra* Quong – a person can forfeit a moral claim she would otherwise have against others *without* failing to treat others in accordance with the moral claims they in fact exert against her. She can do so by taking a morally optional risk.

Moreover, as long as one concedes some intuitive force to the implication of the luck-egalitarian principle that I have just stated (and therefore to the luck-egalitarian principle), one has sufficient reason to reject Quong's proposal in favour of my alternative. The relevant egalitarian intuition says that there is something unfair about the (full) costs of one person's choices being offloaded onto an innocent 3rd party. The parallel intuition, in the context of preventive justice, says that there is likewise something unfair about the (full) harm resulting from one person's risk-imposing activity being offloaded onto an innocent 3rd party. In immunizing the driver against liability, Quong's proposal contravenes this intuition. We might say that Quong sacrifices

distributive fairness for a particular, demanding, conception of deontic fairness.

I imagine Quong would respond to this challenge by reiterating his belief that considerations of distributive fairness (at least distributive fairness in the luck-egalitarian sense) are irrelevant to questions of liability to defensive harm. But the only non-theory-driven reason Quong offers to support this belief is that considerations of distributive fairness yield the counter-intuitive judgment that the driver is liable to be lethally targeted with a grenade. I will argue, however, that such considerations do not support that judgment. Rather they support the weaker (and I think independently plausible) judgment that the driver is liable to be defensively killed only as a side effect.

In summary, at least two different intuitions about fairness are pertinent in the driver case. The first intuition tracks the notion of fair treatment. The second tracks the notion of distributive fairness. The concern with fair treatment directs our attention to the constraints operating on how the driver may be treated, given the way she treats others. The concern with distributive fairness directs our attention to the “innocent other” who is “unfortunately” threatened by the driver’s actions, and to what costs it is fair to allow to be imposed on her as a result of the driver’s actions. Quong’s proposal offers a *prima facie* plausible articulation of the former, but it does so at the cost of the latter. By contrast, the proposal I offered reconciles these intuitions about distributive and deontic fairness.

I’ll conclude this section with an example, originally Otsuka’s, to induce support for the claim that distributive fairness matters for liability. Otsuka contrasts the conscientious driver case with the following variation. Once again, an out-of-control car is about to hit a pedestrian. But this time there is no one driving the car, only a helpless

child in the passenger seat. Otsuka contends that it is morally *worse* for the pedestrian to throw the grenade at the car containing only a helpless passenger than to throw it at the car containing the conscientious driver. As he explains it, “the driver becomes liable [to at least some degree of defensive harm] by virtue of [her] moral responsibility for driving a foreseeably potentially lethal vehicle, whereas there is no plausible account according to which the passenger could become liable.”²⁵

III

i.

The key insight creditable to Quong’s moral status account says that it matters for a threatener’s liability whether she wrongs her victim. Otsuka acknowledges this insight, and claims the fact that the driver does not wrong the pedestrian “mitigates the extent” of her liability.²⁶ This claim is propitious, but essentially promissory: it leaves unspecified both the way in which, and the extent to which the driver’s liability is mitigated. Both of these beg specification. So far, I have defended the responsibility account by defending its commitment to distributive fairness as requisite of any plausible theory of liability to defensive harm. That is, I have defended the foundational normative commitment of the responsibility account. But this defense is incomplete. To fully defend the account, I must demonstrate not only that it is founded on plausible normative commitments, but also that, in ascribing liability in particular cases, it can adjudicate among those commitments in a way that is intuitively plausible, and is not *ad hoc*.

²⁵ Otsuka (2016), p. 67.

²⁶ Otsuka (2016), pp. 66-67.

Section II dismantled Quong's moral status account of liability; but it left substantively unaddressed a distinct critical challenge to the responsibility account, which emerged from his offense. This asserts that, because the driver does not wrong the pedestrian, the conclusion that she's liable to be lethally targeted with a grenade is intuitively unacceptable. I agree with this assertion, and accept it as desiderate of my defense of the responsibility account that I show it not to support the stated conclusion.

There are two candidate specifications of the claim that a person's liability is mitigated. First, one might say the maximal harm to which she is potentially liable is reduced. Second, one might say the harms that others must allow themselves to suffer, in order to avoid inflicting on her the maximal harm, are increased. Both of these specifications have merit, and I believe both apply in the driver case. The task of the present subsection is to argue, via the former specification, that, by the lights of the responsibility account, the driver is not liable to the grenade. The successive task, of III.ii, will be to identify and resolve an apparent tension between the two specifications, and in doing so, provide deeper moral argument for the latter. But let me begin, here, by addressing the latter in a preliminary way.

Suppose, in claiming that the driver's liability is mitigated, Otsuka is claiming that others are obliged to bear increased burdens in order to reduce the defensive harm to the driver. This is to reassert something canonical to the responsibility account. The driver's responsibility is substantially reduced by the fact that she blamelessly threatens to harm the pedestrian accidentally (we can now add: without wronging her); this reduction in the driver's responsibility correlates with an increase in the harms others must absorb to avoid killing her.

McMahan fortifies this position, but also alludes to its principal limitation, with the following remarks:

In most situations in which it is unavoidable that some person or persons must be harmed, or perhaps in all such situations, there is in principle some distribution of unavoidable harm among potential victims that would be ideally just. In the ideally just distribution, each person takes his or her fair share of the harm. A liability justification for harming allocates unavoidable harms in the way that best approximates the ideally just distribution.²⁷

An ideally just distribution of ineliminable harm is sensitive not only to each distributant's absolute degree of responsibility for the harm, but also, crucially, to the incremental differences in responsibility among the distributants.²⁸ Since the driver's absolute degree of responsibility is only slightly higher than the pedestrian's – she is minimally responsible, and the pedestrian non-responsible – in an ideally just distribution, the driver's share of the threatened harm would be only slightly larger than the pedestrian's. If the harm were divisible, then the harm to which the driver would be liable would be reduced, accordingly.

The problem is that specifying the driver's liability as mitigated in this way leaves unparried the criticism that the responsibility account yields counterintuitive results in the original driver case. In the original case, the harm is indivisible: it can fall entirely on the

²⁷ Jeff McMahan, "Duty, Obedience, Desert, and Proportionality in War: A Response" *Ethics* Vol. 22 No. 1 (2011), p. 155.

²⁸ McMahan argues the point differently: he says the ideally just distribution is sensitive to *proportionate* differences in responsibility among the distributants. If you're 90% responsible for the fact that a harm must be distributed between us, and I'm 10% responsible for that fact, an ideally just distribution would inflict on you 90%, and on me 10%, of that harm. But McMahan's argument resists exportation to the driver case. McMahan stipulates the pedestrian's responsibility is "0"; so her fair share of the harm would always be 0.

pedestrian, or entirely on the driver. Since the driver is *more* responsible than the pedestrian, the nearest approximation of the ideally just distribution is to distribute the harm entirely onto the driver.²⁹ The driver would therefore be liable to suffer the full burden of harm under the relevant circumstances.

Otsuka's final remarks about the case accord with this conclusion; he concedes that the driver "might well be" liable to having a grenade thrown at her in the pedestrian's defense.³⁰ But this conclusion contravenes the powerful criticism Otsuka purports to rebut, namely, that the responsibility account cannot accommodate the judgment that, since the driver does not wrong the pedestrian, she does not make herself liable to be killed in this way.

The upshot of this discussion is that, in order to defeat the critical challenge, I must argue that the maximal harm to which the driver is potentially liable, or the ceiling on her potential liability, is lowered.

This strategy appears unpromising. It is a precept of the responsibility account that, in cases of indivisible harm, a threatener can be liable to at least as great a harm as that with which she threatens her victim. Since the driver threatens the pedestrian with lethal harm, she is liable to be lethally harmed herself, if that is necessary to save the pedestrian.

The prospects of this strategy improve, however, if one considers not the maximal degree of defensive harm to which the driver is potentially liable, but the separate question of *the ways in which the driver is liable to be treated* in the course of being defensively harmed to that degree. I think the driver's liability is indeed limited in this

²⁹ McMahan (2011), p. 155.

³⁰ Otsuka (2016), p. 68.

respect. In particular (and as teased in the previous section), I think the driver is not liable to be killed as means of saving the pedestrian, but only as a side effect of saving the pedestrian.³¹ If, for example, you show up on the scene in an armored truck, the car-driver would not be liable to your running her off the road. But the car-driver would be liable to your interposing yourself between her and the pedestrian, even knowing the force of the impact will be fatal to her.

Before I explain why I think this view is yielded by the responsibility account, here are two quick observations regarding its intuitive plausibility. First, the claim that the driver's liability is limited with respect to how she's liable to be treated in the course of being fatally harmed correlates with the claim that it is because of *how the driver treats others* – in particular, because she doesn't wrong anyone through her risk-imposing activity – that her liability is mitigated. Second, it offers a simple way to resolve the apparent tension between the following two claims, both of which are hard to reject: first, that the maximum amount of defensive harm to which a person is potentially liable is at least the amount of harm with which she threatens another, and second, that the driver is not liable to be targeted with a grenade (or run off the road).

Let me now argue that the driver is only liable to be defensively killed as a side effect by the lights of the responsibility account. According to the account, choosing to engage in a foreseeably risk-imposing activity, the risk of which eventuates in a threat of unjustified harm to another person,³² is the basis of liability to defensive harm. If the risk-imposer contravenes one of that other person's rights, the extent of her liability will increase in proportion to the stringency of that right, and in proportion to her degree of

³¹ But see note 33, below.

³² More precisely, to another person who bears a lesser degree of responsibility than the risk-imposer for the fact that one of them will be harmed. (See the example in note 11, above).

blameworthiness for its contravention. Since the driver does not contravene anyone's rights, the ceiling on her potential liability must be fixed in luck-egalitarian terms.

The luck-egalitarian principle, applied in the domain of preventive justice, says that it would be unfair to redistribute the harms resulting from one person's risk-imposing activity onto others. It says furthermore that if the harms from one person's risk-imposing activity would otherwise befall another, innocent, person, it is not unfair to redistribute those harms back onto the risk-imposer.

On the face of it, it looks as though, when the pedestrian targets the driver with a grenade, she's acting in accordance with the luck-egalitarian principle: she is redistributing the harm of the driver's risk-imposing activity back onto the driver, to avoid bearing the full, fatal brunt of it herself. But there is an asymmetry between the claim that the driver is liable to have the harm resulting from her risk-imposing activity redistributed onto her in order to save the pedestrian and the claim that the driver is liable *to be lethally targeted* to save the pedestrian. The luck-egalitarian principle limits the burdens others can be asked to bear on one's behalf. But it doesn't follow from this that people have *carte blanche* to avoid those burdens *by any means necessary*. Accordingly, it does not follow that the driver is liable to *having a grenade thrown at her* as a means of saving the pedestrian. More generally, it does not follow that the driver is liable to be intentionally killed as a means of saving the pedestrian.³³

³³ I'm using 'killed as a means' in the sense of 'killed as a means to an end,' or killed eliminatively. There's a second sense of 'killed as a means,' roughly, 'used as a means, and thereby killed,' or killed opportunistically. Opportunistic killings are thought to be more difficult to justify than eliminative killings; they are outside the concern of this paper. For the original discussion of the distinction between opportunistic and eliminative killings, see Warren Quinn "Actions, Intentions, and Consequences: The Doctrine of Double Effect," *Philosophy & Public Affairs* Vol. 18 No. 4 (1989), pp. 334-51.

I think the luck-egalitarian principle supports a more limited ascription of liability to the driver. In particular, I propose that the driver is liable to be killed not as a means, but only as a side effect of saving the pedestrian.

An interlocutor might accept my claims that (1) there are reasons of fairness to redistribute the harms of the driver's risk-imposing activity back onto the driver, and (2) these reasons do not give would-be redistributors moral license to do so by any means necessary. But my interlocutor might question my move from these claims to the proposal I have advanced, namely, that reasons of fairness are sufficient to make it the case that the driver is liable to be killed as a side effect.

My defense of this move is grounded in the following two, seemingly unrelated observations. First, remember that, if a person is liable to a particular measure of defensive harm, it means she has *forfeited her right* against being harmed in that way, in that circumstance. In the absence of such forfeiture, she would *possess* a right against being harmed in that way. So the claim that the driver is liable to having the harms of her risk-imposing activity redistributed onto her in some way entails that she has done something to forfeit a right she would otherwise possess against having the harms in

The view that the pedestrian kills the driver as a means of saving herself is controvertible. The opposing view is that the pedestrian *blows up the driver's car* as a means of saving herself, and only kills the driver as a foreseen side effect of blowing up the car. In deference to the opposing view, I elsewhere argue, using Frances Kamm's account of subordination, that (1) certain side-effect killings are relevantly like means killings in the respect that they subordinate their victims, (2) these side-effect killings are more difficult to justify than side-effect killings that do not, (3) throwing a grenade at the car kills the driver as a side effect in a way that subordinates her, and (4) interposing a shield between the pedestrian and the car does not. ("Not as a Means: Killing as a Side Effect in Self-Defense," unpublished manuscript.)

From the point of view of the argument of the present paper, not a lot hangs on whether, by throwing a grenade at the car, the pedestrian kills the driver as a means, or as a side effect in a way that subordinates her. I'll address the main implication for my argument here if the latter view is correct in note 35, below.

For defense of the former view, see Helen Frowe *Defensive Killing* (Oxford: Oxford University Press, 2014), pp. 38-39.

question redistributed onto her in that way.

Second, recall the non-consequentialist ranking of different types of killings, which identifies killing someone as a means as morally worse than killing someone as a side effect,³⁴ and thereby identifies a person's right against being killed as a side effect as less stringent than her right against being killed as a means. Consider now that there is a further non-consequentialist distinction, which identifies being *killed* as a side effect as morally worse than being *allowed to die* as a side effect. Non-consequentialists obviously accept that there's moral reason not to allow someone to die as the foreseen side effect of one's action. But they deny that there's a general *right* against being allowed to die as a side effect. Thus, for example, in the case in which a car containing only a helpless passenger threatens a pedestrian with fatal harm, the pedestrian would be permitted to take evasive action (to leap out of the path of the car, say), even foreseeing the car would subsequently crash into a building, killing the passenger inside. It's not that the passenger has forfeited her right not to be allowed to die in this circumstance – the passenger is completely nonliable, which means she has not forfeited any of her rights – it's that the passenger never had such a right in the first place.

My claim that the luck-egalitarian principle supports the view that the driver is liable to be killed as a side effect of saving the pedestrian is thus supported by the following:

(1) For the claim that the driver is liable to have the harms of her risk-imposing activity redistributed back onto her in some way to be meaningful, it must be the case that she has *forfeited a right*, which she would otherwise possess, against having the harms of her risk-imposing activity redistributed onto her in that way.

³⁴ There are gradations among different types of means-killings, too; see note 33, above.

(2) Being killed as a side effect and being allowed to die as a side effect are next to each other on the non-consequentialist continuum. That is, there's no way of fatally harming a person that's both morally better than killing her as a foreseen side effect, and morally worse than allowing her to die as a foreseen side effect.

(3) Given that a person lacks a general right against being allowed to die as a side effect, but possesses such a right against being killed as a side effect, the right against being killed as a side effect is the least stringent general right a person has against fatal harm.

(4) If a person is liable to have a fatal harm redistributed onto her, she must be liable at least to be killed as a side effect, and in the absence of other moral considerations (i.e., blameworthiness) serving to increase her liability, there's no reason to believe that the extent of her liability goes beyond the liability to being killed as a side effect.³⁵

The responsibility account thus yields the view that the ceiling on the driver's liability to defensive harm is fixed at liability to be killed as a side effect. This is the maximal liability that is supported by the luck-egalitarian principle; and the driver's liability is not increased by other considerations.

³⁵ The "cut" I make in (2), and apply in (3), between killing a person as a means and killing a person as a side effect, may be too coarse grained. I argue in "Not as a Means" (and synopsise in note 33, above) that a further cut can be made between two types of side-effect killings, namely, those that subordinate their victims and those that do not. The former are morally worse than the latter; accordingly, a person's right against the former is more stringent than her right against the latter.

The implication, for the present argument, of this being correct, would be to strengthen the argument's conclusion, by further restricting the ways in which the driver is potentially liable to be killed. The argument dictates the driver only forfeits the least stringent general right a person has against being killed. Appropriately amended, the argument would entail not that the driver forfeits her right against being killed as a side effect, *simpliciter* (because this right is too coarsely demarcated), but only that she forfeits her right against being killed as a side effect in a way that does not subordinate her.

One question arising here concerns whether the luck-egalitarian principle supports the view that the driver is potentially liable to be defensively harmed to any extent as a means, or supports the view that she is only liable to be harmed as a side effect. I think the most plausible answer is that fixing the ceiling on the driver's liability at being killed as a side effect entails that the driver does not forfeit any rights of *greater stringency* than the right against being killed as a side effect. But she potentially forfeits any rights of stringency equal to, or lesser than, the right against being killed as a side effect. Since that set of rights presumably includes the right against being inflicted with certain lesser harms as a means, she is potentially liable to be harmed in those ways as a means.³⁶

ii.

The discussion of III.i prompts the following objection. I advanced versions of these three claims:

(1) It is unfair if the harms resulting from one person's risk-imposing activity are borne by non-responsible third parties.

(2) It is not unfair if the harms resulting from a person's risk-imposing activity are borne by the risk-imposer.

(3) The fair distribution of harms resulting from the driver's risk-imposing activity between the risk-imposing driver and the non-responsible pedestrian allocates a share of the harm onto each, with the driver's share being only a modest amount larger than the pedestrian's.³⁷

³⁶ I won't seek to specify here how much harm, inflicted as a means, is equivalent to the harm of death, inflicted as a side effect.

³⁷ I assume the driver and the pedestrian are the only candidate distributants for the relevant harm. If the pool of candidate distributants were expanded, the fair distribution of harm would allocate

(1) and (2) undergird my account of the driver's liability. (3) is the explanation I endorsed for the pedestrian's obligation to burden herself in order to avoid imposing on the driver the maximal harm to which the driver is potentially liable. (That is, it's the explanation I endorsed for the driver's liability being mitigated in the first sense discussed in III.i). But, of course, (1) and (2) together contradict (3). So, the objection says, I need to revise or abandon either my account of the driver's liability, or my account of why others are obliged to substantially burden themselves in order to avoid inflicting on the driver the maximum harm to which she is potentially liable.

The claim advanced in (3) strikes me as correct. I think, moreover, that it reveals a certain crudeness in the formulations of (1) and (2), and thus, in my account of the driver's liability. So let me address the objection by articulating and defending a refined version of my account of the driver's liability, starting with refined versions of (1) and (2):

(Refined 1) It is unfair if the *full* harms resulting from one person's risk-imposing activity are borne by non-responsible third parties.

(Refined 2) It is not unfair if *the largest share of* the harms resulting from a person's risk-imposing activity is borne by the risk-imposer.

These refined versions of (1) and (2) accord with (3) (above): they accord with the claim that fairness prefers the driver and the pedestrian each to absorb a share of the harm, and allows for the difference between their respective shares to be modest.

Furthermore, (Refined 1) and (Refined 2) together articulate a superior conception of

smaller shares onto more people. But the driver's share would presumptively remain larger than anyone else's.

luck-egalitarian preventive justice than the original versions of (1) and (2) do. This is because, unlike the original versions, they reflect contemporary advancements in the specification of the luck-egalitarian ideal. In particular, they reflect the more nuanced way contemporary egalitarians conceive of the proper role of choice in shaping distributions.

On the way in which choice influences, and does *not* influence, the justice, or fairness, of distributions, G.A. Cohen writes that:

You do not escape responsibility for the costs of your choice by virtue of the mere fact that you made that choice *against a choice-affecting background*. But *also*, the mere fact that you made a choice, and could have chosen otherwise... [does not show] that subsidy is out of order.... [F]acts in the background to the choice, facts about degrees of control, and about the costs of alternatives, affect the proper allocation of responsibility for the consequences of that choice.³⁸

The fact that a person made a choice that resulted in her becoming worse off than others generally *reduces* the extent of her fairness-based claim to compensation, relative to what it would be if her loss were entirely the product of bad brute luck. But the mere fact that she made such a choice is not sufficient to make it the case that she has *no* fairness-based claim to compensation at all.

Cohen's point – that the fact that a person made a choice does not settle the question of the extent to which it is appropriate to hold her responsible for bearing its

³⁸ G.A. Cohen, "Expensive Taste Rides Again," first published in 2004, reprinted in *On the Currency of Egalitarian Justice, and Other Essays in Political Philosophy*, Otsuka ed., (Princeton: Princeton University Press, 2011), p. 109. (Italics preserved from the original). See also Cohen, "On the Currency of Egalitarian Justice," first published in 1989, reprinted in Otsuka ed., p. 32.

costs – can be transposed into the domain of preventive justice. The mere fact that a person chose to engage in a risk-imposing activity, the risk of which will eventuate in harms of a fixed magnitude, does not settle the question of *the extent to which* she forfeits her claim against others to shoulder some of those harms on her behalf.³⁹ “Facts in the background to the choice” will also be pertinent to this question.

Before I address which – if any – background facts mitigate the extent to which the driver can appropriately be held responsible for bearing the costs of her choice (and thus increase the extent to which others are obliged to absorb some of those costs), I should make explicit that the present discussion is tangential to, and so does not prompt revisiting, the question of the *ceiling* on the driver’s liability. *That* question arises with respect to indivisible harms. And since the *fact* that a person made a choice always bears on the question of the fair distribution of ineliminable harms, such that, in a perfectly fair distribution, the ‘chooser’s’ share will be the largest, considerations of fairness will still recommend that the driver bear the full harms of her choice just in case those harms are indivisible. As stated in III.i, following McMahan, if one person must bear the harms, the nearest approximation of the ideally fair distribution allocates that harm to the driver.

I think certain facts in the background to the driver’s choice to engage in a risk-imposing activity indeed limit the extent to which she forfeits her claim against others to suffer some of the resultant harms on her behalf. Let me consider two, the first of which pertains to the *costliness*, or *burdensomeness*, of the alternatives available to the driver.

It’s a truism that the organization of our society is predicated on people relying on

³⁹ McMahan says as much; he identifies blame, excuses and non-consequentialist considerations among the factors influencing (by either mitigating or exacerbating) the extent of a person’s liability. The present discussion supplements McMahan’s account by identifying additional mitigating factors.

driving in conducting their day-to-day lives. (Even non-drivers rely extensively on the driving of others, directly and indirectly). In this context, eschewing driving in carrying out one's daily routine (including pursuing certain leisure activities) can be extremely costly. In virtue of this fact, the choice to drive in the course of pursuing an ordinary activity, such as going to the cinema, is disanalogous with paradigm examples of choices the costs of which egalitarians would be disinclined to subsidize, such as the choice to invest some of one's disposable income in a particular stock. I submit that the driver can properly be held responsible for the costs of her choice to a lesser extent than the stock-investor can, and thus, that the costs others can be asked to bear on her behalf are correspondingly higher.

The second background fact is that driving – or more precisely, driving in accordance with certain regulations and norms – is a permissible type of activity. It will of course be the case that, on virtually every occasion on which a person drives, she will expose herself to, and impose on others, at least a slight risk of serious harm. But it is generally accepted as given that the risk of harm one person's careful driving imposes on others is tiny – “so tiny that the activity, considered as a *type* of [risk-imposing] activity, is entirely permissible.”⁴⁰ I believe – and for the purposes of this paper, will assume – that the justification for classifying driving as a permissible type of risk-imposing activity has the following structure: the collective benefits we accrue (and I assume we each benefit, even if we do so unequally) from the practice of driving (or, again, of driving in accordance with certain regulations and norms) gives us each decisive reason to accept

⁴⁰ McMahan (2009), p. 165.

the small risk to which that practice exposes us.⁴¹ In other words, the benefits of living in a society in which driving is normalized (in roughly the way in which driving is normalized in ours) are, to each member, worth the risk to which she will be exposed by the (careful, law-abiding) driving of others.

It's worth making explicit the significance of the formulation that driving is a permissible type of activity *just in case it's undertaken in accordance with certain regulations and norms*. First, and most straightforwardly, the riskiness of driving is kept at acceptably low levels (that is, at levels low enough to be justifiable to each person exposed to that risk) by strict regulations about speed limits, traffic signals, auto-safety standards and so forth, and by norms of drivers further 'self-regulating' their speed, etc., when driving in particularly harsh weather conditions, for example. But there is a second, less obvious way in which the set of rules and norms that are in effect will determine, in part, the riskiness of driving, which is by regulating the distribution of the harms concomitant with permitting driving as a ubiquitous social practice. In other words, the risk to which a person exposes herself (and imposes on others) by driving will be, in part, a function of the rules and norms governing the distribution of its resultant harms.

Take the following two alternative principles for distributing the harms of driving in cases like the conscientious driver's. The first, which I'll call the principle of strict liability, says that the driver is liable to have redistributed onto her the full harm that would otherwise fall on an innocent 3rd party as the result of her driving, regardless of

⁴¹ I owe my characterization of driving's permissibility, and much of the subsequent discussion of its implications, to Rahul Kumar, "Risking and Wronging," *Philosophy and Public Affairs*, Vol. 43 No. 1 (2015), (unpublished draft August 2013). This discussion is also indebted to Quong (2012) pp. 63-65.

whether that harm is divisible. The second says that the driver is liable to have redistributed onto her some percentage of the harm that would otherwise fall on an innocent 3rd party as the result of her driving, but less than 100% of that harm just in case the harm is divisible. If a principle of strict liability is adopted, the number of victims per car accident will be reduced, and so fewer people will be harmed in car accidents over time. Thus, the *risk* each person will face of being harmed in a car accident will be reduced (at least relative to what it would be if something like the 2nd principle is adopted). But it will also be the case, if strict liability is adopted, that the severity of the harms suffered by the victims of car accidents will be increased, overall. Assuming that the magnitude of the harm resulting from each accident is fixed, smaller numbers of victims will entail a larger amount of harm to each victim. In this respect, the risk each person will face that *if* she's harmed in a car accident, she'll be harmed severely, will be increased under a principle of strict-liability (again, relative to under the second principle). Each of the two candidate principles thus makes a different trade-off between the risk to each person of becoming a victim, and the risk to each victim of being severely harmed.

A person might reasonably be indifferent between a lesser risk of suffering a greater harm, and a greater risk of suffering a lesser harm. But what is in persons' interests, narrowly construed, is not the only relevant factor for selecting between the two alternatives. Considerations of the values in accordance with which we have most reason to prefer society to be governed also matter. And here, I think the value of mutuality (or community) recommends decisively against a principle of strict liability for the harms that result from permissible risk-imposing activities. It recommends, instead, that

individuals (like the pedestrian) absorb some of the harms resulting from others' permissible risk-imposing activity, if doing so is necessary to avoid redistributing the full harm of that activity back on the risk-imposer.^{42, 43}

The opposition between mutuality and strict liability is limited to cases involving *permissible* risk-imposing activities. This is because that value relies on certain norms of reciprocity being in effect. Consider: from the point of view of mutuality, society is conceived, in the first instance, as a cooperative endeavor. When a person acts in accordance with the value of mutuality, she demonstrates a willingness to absorb certain costs or burdens on others' behalf (in ways that are distinct from, and go beyond, the luck-egalitarian commitment to equalize the burdens of bad brute luck). But this orientation can become a form of servility if it is one-sided, that is, if people are not *mutually* willing to absorb costs on one another's behalf. When a person engages in a permissible risk-imposing activity (i.e., when she drives carefully and alertly), she is (on that occasion and to that extent) 'doing her part' to keep the overall level of risk to which her fellow members of society are exposed acceptably low. She thereby acts in accordance with norms of reciprocity. By contrast, when a person engages in an *impermissible* risk-imposing activity (for example, when she drives recklessly to the cinema to avoid missing the start of a film), she flouts those norms. In this respect, she identifies herself as a 'non-cooperator,' and thereby exempts others (again, on that occasion and to that extent) from the mutuality-based obligations they would otherwise have toward her.

⁴² Or if doing so is necessary to avoid the full harm of that activity being redistributed back onto the risk-imposer by a third party, on one's behalf.

⁴³ The discussion in this paragraph and the next is indebted to Seana Shiffrin "Paternalism, Unconscionability Doctrine, and Accommodation," *Philosophy and Public Affairs* Vol. 29 No. 3 (Summer, 2000), Section II.B, pp. 242-245.

Let me summarize the argument I've just laid out. The level of risk concomitant with driving is, in part, a function of the rules and norms that govern the practice, as well as of the rules and norms that govern the distribution of its ensuing harms. Driving is a permissible type of risk-imposing activity because its governing rules and norms keep the risk of harm to which it exposes each person sufficiently low, such that it is in each person's interest to prefer living in a society in which driving (in accordance with those rules and norms) is permitted to living in a society in which driving is prohibited. In this respect, the permissibility of driving is fixed with regard to the way in which people have most reason to want their society to be organized. But the question "how do people have most reason to want their society to be organized?" is pertinent not only to the selection of the rules and norms for governing the practice of driving, but also to the selection of the rules and norms for distributing the harms that result from driving. Granted, individuals might be indifferent between exposure to a lower risk of suffering a more severe harm, and a higher risk of suffering a less severe harm. But the scope of the question "how do people have most reason to want their society to be organized?" is not limited to what's in an individual's interest, narrowly construed; rather, it extends to include the question "what are the *values* in accordance with which people have most reason to want their society to be organized?" I contended that one of these values is the value of mutuality. And, from the point of view of mutuality, there is decisive reason to prefer distributing the harms associated with permissible driving in a way that exposes each person to a greater risk of suffering less severe harms.

The upshot of this discussion is as follows. The set of considerations that determine the extent to which it is appropriate to hold the driver responsible for bearing

the harms that result from her driving includes the set of rules and norms that *ought* to govern, or the correct set of rules and norms for governing, the distribution of harms resulting from driving. Since the driver is conducting herself in accordance with that set of rules and norms – and thereby conducting herself in accordance with the norms of reciprocity – she has a claim against others to treat her in accordance with the value of mutuality. In particular, she has a claim against others (and others, correlatively, have an obligation toward her) to accept a share of the divisible harms that result from her driving, rather than redistributing 100% of those harms back onto her.

The significance of the *permissibility* of driving (or, more precisely, the significance of the correct set of rules and norms for governing the distribution of harms that result from permissible driving) to the extent of the driver's liability is different from the significance of the *costliness* of eschewing driving to the extent of the driver's liability. The costliness of eschewing driving makes it distributively unfair to hold the driver responsible for the full cost of that choice. By contrast, the permissibility of driving reduces the driver's liability by tempering the demands of distributive fairness.

Let me explain. Distributive fairness, on the luck-egalitarian conception, is concerned with the extent to which persons' choices are reflected in the costs to which they are subject: it commands an equal distribution of costs insofar as they do not reflect choice. The degree to which someone's plight reflects her choices will be a function of several things, including "the character of the alternatives the chooser had."⁴⁴ So, for example, when a person chooses a particular option against a background in which all of the alternatives were too expensive to be viable for her, the outcome reflects her choice less (in the relevant sense) than it would if she were choosing among a set of viable

⁴⁴ Cohen, p. 108.

alternatives. Judgments of distributive fairness will thus be sensitive to (considerations such as) the costs of the alternatives that were available to the choosers.

The extent to which a person's choices are reflected in the costs to which she is subject will also be affected by the *actual* set of rules and norms in place for governing the distribution of the harms that may result from the various 'risky' alternatives available to her. This is because the greater the extent to which she expects to have to shoulder the burdens associated with a given risk, the more 'expensive' she will perceive that risk to be. I think the potential harms associated with certain risks can have choice-constraining (or choice-distorting) effects,⁴⁵ which are roughly analogous to the choice-constraining effects of the advance costs of various options.⁴⁶ Judgments of distributive fairness will thus be sensitive (to at least some degree) to the rules and norms in place for governing the distribution of the burdens that arise when people's risks eventuate.

By contrast, the extent to which a person's choices are reflected in the costs to which she is subject will *not* be contingent on what the *correct* set of rules and norms is for governing the distribution of costs or harms that result from her optional risk-taking. Normative considerations of this kind do not factor into people's practical deliberations in the way that the actual costs of, or risks associated with, their various options do. Judgments of distributive fairness will therefore be *insensitive* to the correct set of rules and norms for governing the distribution of burdens resulting from people's choices.

The upshot of this discussion is that we should reject my earlier assumption that the question "to what extent are people obliged to absorb some of the costs (or harms) of

⁴⁵ For defense, see my "Egalitarianism for Girls" (working title), unpublished manuscript.

⁴⁶ One might suggest that a policy of strict liability to the (unavoidable but potentially divisible) harms that result from permissible driving can itself constrain choice in this way. But I think the risk in question is too tiny to constrain the choice of persons whose psychology is within the normal range.

someone's else's choices?" is equivalent – or reducible – to the question “what is the extent of a person's fairness-based claim to compensation for the costs of her choices?”

The foregoing argument reveals these questions to be distinct from one another.

That argument also reveals what I believe to be a mistake in the way the luck-egalitarian conception of distributive justice is pervasively demarcated, by both its proponents (including Cohen), and its critics.⁴⁷ The scope of luck-egalitarian justice is construed as co-extensive with that of distributive fairness, which is distinguished by its impulse to equalize all, and only, those costs that do not result from choice. But the scope of luck-egalitarian justice more plausibly correlates with the scope of the question “to what extent can a person appropriately, or properly, be held responsible for bearing the costs that result from her choices?” That is, the set of normative considerations to which luck-egalitarian justice is sensitive includes, but is not limited to, those of distributive fairness.⁴⁸ This proposal is supported by Cohen's claim that luck-egalitarian justice is concerned with the “proper” or “reasonable” allocation of responsibility for the consequences (i.e., for the harms or costs) of persons' choices, on the reading of that claim for which I have argued. It is of course an open question whether Cohen would accept this reading of ‘proper’ and ‘reasonable’ (according to which they track certain considerations outside the set that he explicitly raises), or affirm the narrower reading (according to which they track a more limited set of considerations, which does *not* include the correct set of rules and norms for governing the distribution of burdens arising from people's choices), and thus, reaffirm the standard conception of luck-egalitarian justice. But the main claims I want to advance, and for which I have argued

⁴⁷ I argue for this claim from a different direction in “Luck, Love, and Extreme Skiing: Distributive Injustice without Unfairness,” *Moral Philosophy and Politics*, Vol. 3 No. 1 (2016).

⁴⁸ If this is correct, then “luck egalitarianism” is a misnomer, but never mind.

are (1) that the broader conception of egalitarian justice that I am proposing has greater intuitive plausibility than the standard luck-egalitarian conception and (2) that it is rooted in considerations contiguous with those already picked out, by Cohen, as internal to the luck-egalitarian ideal.

In this subsection, I have argued, first, that the fact that a person *chose* to engage in a risk-imposing activity is always pertinent to the question of how to distribute the harms that result when the risk concomitant with the activity eventuates. In particular, it's sufficient to establish that, under an ideally just distribution, the *largest* share of the harm will be borne by the risk-imposer. But the fact that a person chose to engage in a risk-imposing activity, the risk of which will eventuate, does not settle the question of how *much* larger her fair share – broadly construed as “ideally just” share – will be than those of others. The answer to that question will depend, in part, on facts in the background to her choice. With respect to the conscientious driver, I identified two relevant background facts, each of which, I argued, reduces the driver's fair share of the harm in question. The first of these pertains to the costliness, or burdensomeness of the alternatives available to her. The second pertains to the permissibility of driving in the manner (and in the societal context) in which she drives. I submit that these two facts work in tandem to reduce the difference between the driver's fair share of the harm, and other peoples' fair shares of the harm, to a small amount.

IV

The task of section III was positive; it advanced a case for how I think the implications of the responsibility account should be cashed out, using the conscientious driver as the motivating example. I presented the arguments in that section as developing features already implicit within McMahan's articulation of the account. But the version of the account I advanced deviates from McMahan's. McMahan identifies the driver as blameless (or non-culpable) for the threat she poses on the grounds that she is excused; he thereby identifies the driver as doing something morally wrong.⁴⁹ By contrast, my claim that the driver is liable to defensive harm on luck-egalitarian grounds does not appeal to the modality of rightness or wrongness. The sole moral consideration appealed to is fairness. The view says simply that the fact that the driver chose to engage in a risk-imposing activity that will now have fatal consequences makes it fair to allocate onto the driver the largest share of the resultant costs or harms.

That McMahan thinks wrongdoing is a necessary condition for liability is implicit in his discussion of blame and excuses. According to McMahan (and as reconstructed in I.iii), if a person is liable to defensive harm, she is either blameworthy or excused. Her liability increases to the extent that she is blameworthy, and decreases to the extent that she is excused. Blame and its exculpation both imply wrongdoing: a person only becomes a candidate for moral blame or its exculpation by doing something morally wrong.⁵⁰ But this claim does no load-bearing work in McMahan's account of liability. A person becomes liable to defensive harm when she chooses to engage in a risk-imposing

⁴⁹ McMahan (2009), p. 165.

⁵⁰ McMahan (2009), p. 112.

activity, the risk of which will eventuate in a threat of unjustified harm to another. Adding that *if* a person engages in a risk imposing activity, the risk of which will eventuate in an unjustified harm, then she's doing something morally wrong, does no explanatory or justificatory work in establishing the fact (as opposed to the extent) of her liability.⁵¹

If the claim that liability implies wrongdoing were a freestanding assumption of McMahan's, it could be dismissed. But it is a direct implication of one of his express theoretical commitments. McMahan endorses, within the context of ascribing liability, a "fact-relative" account of permissibility,⁵² according to which the permissibility of an act is fixed in relation to facts independent of the agent's beliefs about its permissibility.⁵³ According to the fact-relative account, an act is permissible just in case it *would be* permissible if all the facts were as the agent believed them to be; it is impermissible just in case it *would be* impermissible if the agent knew all the relevant facts.⁵⁴ For example, in the driver case, the driver (justifiably) believes that the risk of unjustified harm she imposes on others will not eventuate. If her belief were true, she would be acting in a way that is fact-relative permissible. But in fact, the risk of unjustified harm she imposes

⁵¹ I suspect McMahan has simply inherited the assumption that wrongdoing is a necessary condition for liability from the culpability account of liability, which he opposes. According to the culpability account, blameworthiness (or culpability) is a necessary condition for liability. McMahan argues, against the culpability account, that blameworthiness isn't a necessary condition for liability because it isn't a necessary condition for *wrongdoing*: a person, such as the driver, can act wrongly, but with a full exculpating excuse ([2009], p. 112). A person who acts wrongly but with a full excuse is nevertheless potentially liable. McMahan rejects the culpability account's contention that blameworthiness is a necessary condition for liability, while retaining its assumption that wrongdoing is a necessary condition for liability.

⁵² Derek Parfit, *On What Matters*, Vol. 1 (Oxford: Oxford University Press, 2011), pp. 150-151.

⁵³ McMahan (2009), p. 43. McMahan's term, "objective permissibility," is synonymous with "fact-relative permissibility."

⁵⁴ I've borrowed this formulation from Quong (2012), p. 48.

on others *will* eventuate and so she is acting in a way that's fact-relative impermissible.⁵⁵ (Since this fact cannot be known *ex ante*, McMahan says she is fully excused or exculpated of blame).⁵⁶ The claim that liability implies wrongdoing might thus be interpreted as a spandrel of McMahan's account: a person becomes liable to defensive harm by engaging in a risk-imposing activity, the risk of which will eventuate in an unjustified harm; and the *fact* that a risk will eventuate in an unjustified harm is a sufficient condition for it being impermissible to impose.

There's a good reason, however, to reject the fact-relative account of permissibility as a load-bearing feature of the responsibility account, which is that it has implausible implications with respect to the permissibility of risk-imposition. According to the fact-relative account, if a risk of an unjustified harm will eventuate, then that risk is impermissible to impose on others. In other words, the fact that a risk imposed on others will eventuate in an unjustified harm is a *sufficient condition* for its impermissibility. Correlatively, that a risk imposed on others will *not* eventuate in an unjustified harm is a necessary condition for its permissibility.⁵⁷ The requirement that a risk must not eventuate in order for its imposition to be permissible is deceptively banal. First, it contravenes the (broadly contractualist) characterization of permissible risk-imposition I employed in III.ii, which says that a risk-imposing activity (such as driving) is

⁵⁵ McMahan (2009), p. 165.

⁵⁶ McMahan (2009), p. 166.

⁵⁷ I take it while McMahan identifies fact-relative permissibility as a necessary condition of permissible risk-imposition, he does not regard it as a sufficient condition. Thus, he can maintain that an activity like driving recklessly to the movies is an impermissible risk-imposing activity, regardless of whether it eventuates in a threat of unjustified harm. Risk-imposing activities like reckless driving would be fixed as impermissible *ex ante*. By contrast, risk-imposing activities like conscientious driving would be of unknown permissibility *ex ante*; the normative valences of tokens of these types of activities would be fixed *ex post*. (This might be the best explanation of what McMahan has in mind when he says that conscientious driving is a permissible type of activity [2009], p. 165).

permissible if the benefits derived from that activity are, to each person, worth the risk to which it exposes her. It also contravenes the standard consequentialist account of permissible risk-imposition, which indexes the permissibility of a risk-imposing activity to the benefits and burdens expected to accrue from that activity's practice over time. In this respect, identifying fact-relative permissibility as a necessary condition of permissible risk-imposition is a radical proposal (relative to both common intuitions and prevailing philosophical theory); it thus requires argument in its defense, which McMahan does not provide.

The claim that wrongdoing is a necessary condition for liability does no work in the responsibility account; and it is undergirded by an implausible theoretical commitment. The account can be amended without loss as follows. First, familiarly, it's a necessary condition for liability to defensive harm that a person chose to engage in a risk-imposing activity, the risk of which will eventuate in an unjustified harm to another. The risk-imposing activity in question may be permissible or impermissible. When the risk-imposing activity is *impermissible* – for example, when a person is driving recklessly, or carrying out a murderous attack – her liability will increase in proportion to her degree of blameworthiness; it will decrease in proportion to the degree to which she is excused. On this version of what is necessary for liability, it doesn't follow that if a person is liable and blameless, she must be excused. She might be excused, or she might not be doing anything wrong.

CONCLUSION

In this paper, I sought to defend a version of McMahan's responsibility account of liability against Quong's bipartite attack. The positive component of Quong's attack advances a counter-argument to McMahan in support of Quong's own moral status account of liability. The negative component purports that the responsibility account yields an intuitively unacceptable judgment in the driver case (and in other cases like it). My defense of the responsibility account was likewise twofold. First, I rejected Quong's moral status account on the grounds that it relies on an implausible and undefended premise, which is that wronging a person is a necessary condition for forfeiting any of one's rights (or moral claims). Second, I defended the responsibility account against the charge of implausibility by arguing that the luck-egalitarian principle that undergirds it supports a more limited construal of driver's liability than Quong (or McMahan) has recognized. In making this case, I cashed out the implications of the luck-egalitarian principle in the domain of preventive justice. I thereby sought to fill a significant lacuna in the canonical account.⁵⁸

⁵⁸ This paper is indebted to Saba Bazargan, Louis Philippe Hodgson, Will Kymlicka, Andrew Lister, Dave Langlois, Margaret Moore, Michael Otsuka, Thomas Pogge, Jonathan Quong, Henry Richardson, Andrew Ross, Victor Tadros, Aness Webster, Andrew Williams, and to members of the Queen's University Political Philosophy Reading Group, and the University of Toronto Ethics and Social Philosophy Discussion Group. Its deepest debts are to Jeff McMahan, whose detailed notes elevated the paper twice over; and to Rahul Kumar, a tireless interlocutor and editor, who midwived the paper into its present form.

REFERENCES

- Bazargan, Saba 2014. "Killing Minimally Responsible Threats," *Ethics* 125/1: 114-136.
- Cohen, G.A. 2011. *On the Currency of Egalitarian Justice, and Other Essays in Political Philosophy*, ed. Michael Otsuka, Princeton: Princeton University Press.
- Frowe, Helen 2014. *Defensive Killing*, Oxford: Oxford University Press.
- Gordon-Solmon, Kerah. Luck, Love, and Extreme Skiing: Distributive Injustice without Unfairness, *Moral Philosophy and Politics*, 3/1: 119-133.
- Egalitarianism for Girls (working title, unpublished manuscript).
 - Not as a Means: Killings as a Side Effect in Self-Defense (unpublished manuscript).
- Kumar, Rahul 2015. Risking and Wronging, *Philosophy and Public Affairs* 43/1: 27-51.
- McMahan, Jeff 2005. The Basis of Moral Liability to Defensive Killing, *Philosophical Issues* 15: 386-405.
- 2009. *Killing in War*, Oxford: Oxford University Press.
 - 2011. Duty, Obedience, Desert, and Proportionality in War: A Response, *Ethics* 122/1: 135-67.
 - 2014. Self-Defense Against Justified Threateners, in *How We Fight: Issues in Jus in Bello*, eds. Helen Frowe and Gerald Lang, Oxford: Oxford University Press: 104-137.
 - 2016. The Limits of Self-Defense, in *The Ethics of Self-Defense*, eds. Christian Coons and Michael Weber, New York: Oxford University Press: 185-210.
- Otsuka, Michael 1994. Killing the Innocent in Self-Defense, *Philosophy & Public*

Affairs, 23/1: 74-94.

– 2016. The Moral Responsibility Account of Liability to Defensive Killing, in *The Ethics of Self-Defense*, eds. Christian Coons and Michael Weber, New York: Oxford University Press: 51-68.

Parfit, Derek 2011. *On What Matters*, Vol. 1, Oxford: Oxford University Press.

Quinn, Warren 1989. Actions, Intentions, and Consequences: The Doctrine of Double Effect, *Philosophy & Public Affairs*, 18/4: 334-51.

Quong, Jonathan 2012. Liability to Defensive Harm, *Philosophy and Public Affairs*, 40/1: 45-77.

Shiffrin, Seana 2000. Paternalism, Unconscionability Doctrine, and Accommodation, *Philosophy and Public Affairs*, 29/3: 205-250.