



SUSANNE BURRI* 

MORALLY PERMISSIBLE RISK IMPOSITION AND LIABILITY TO DEFENSIVE HARM

(Accepted 25 November 2019)

ABSTRACT. This paper examines whether an agent becomes liable to defensive harm by engaging in a morally permissible but foreseeably risk-imposing activity that subsequently threatens objectively unjustified harm. It first clarifies the notion of a *foreseeably risk-imposing activity* by proposing that an activity should count as foreseeably risk-imposing if an agent may morally permissibly perform it only if she abides by certain duties of care. Those who argue that engaging in such an activity can render an agent liable to defensive harm ground this liability in the luck egalitarian thought that we may justly hold individuals responsible for the consequences of their voluntary choices. Against this, I argue that a luck egalitarian commitment to holding people responsible cannot, by itself, ground liability to defensive harm. It can *help* ground such liability only against the backdrop of a *distributively just society*, and only if *further considerations* speak morally in favour of attaching certain well-defined costs to individuals' risk-imposing choices. I conclude by suggesting that if an account of liability applies robustly across distributively just and unjust contexts alike, then what grounds an agent's liability is plausibly not her responsibility for threatening objectively unjustified harm, but her culpability for doing so.

I. INTRODUCTION

According to the Moral Responsibility Account of Liability to Defensive Harm (Responsibility Account, for short), it can be permissible to harm a

* For helpful discussions and comments, I thank Andreas Brekke Carlsson, Lars Christie, Roberto Fumagalli, Siba Harb, Lisa Hecht, Linda Radzik, Michael Robillard, Thomas Seiler, Henrik Syse, Alex Voorhoeve, the participants at the 8th annual ELAC conference that was hosted in August 2017 in Oslo, Norway, the participants at the March 2018 UCL-LSE Workshop on Risk and Aggregation in Ethics, the participants at the 3rd annual Bovay Workshop on Ethical Risk Management at Texas A&M University in April 2018, as well as an anonymous reviewer for this journal.

person in self-defence even if that person is not at fault for the threat they pose.¹ Consider the example of a *conscientious driver*.² This driver always keeps her car well-maintained and always drives carefully. One fateful evening, her car nevertheless veers out of control, threatening to kill a pedestrian. Defenders of the Responsibility Account claim that if the pedestrian can save his own life only by killing the driver, then he is morally permitted to do so, because the driver is *morally responsible* for threatening the pedestrian, and for this reason lacks a right against whatever necessary and proportionate defensive harm the pedestrian might impose on her. Due to her conscientiousness, the driver is not blameworthy for threatening the pedestrian, but she is nevertheless responsible, as it was foreseeable to her that deciding to drive might result in a threat of unjustified harm. By contrast, if someone threatens unjustified harm through an action where no harm was foreseeable, then they are not responsible for the threat they are posing, and are therefore not liable to defensive harm. Consider the example of an *unsuspecting homeowner* who will cause a lightning flash in a neighbouring house if he flips a light switch in his hallway.³ The homeowner was not able to foresee that flipping the switch might result in a threat of unjustified harm, and is thus neither responsible for threatening his neighbour, nor liable to defensive harm.

In this paper, I critically examine the Responsibility Account. I first explore whether it is possible to make sense of the claim that a person such as the conscientious driver was able to *foresee* that her actions might threaten harm in a way that a person such as the unsuspecting homeowner was not. I conclude that a principled line between the two agents can indeed be drawn. While this is good news for the Responsibility Account, which relies on distinguishing these cases, I go on to argue that the Responsibility Account fails to be backed up by a satisfactory rationale.

¹ Central defences of the Responsibility Account include Jeff McMahan, 'The Basis of Moral Liability to Defensive Killing', *Philosophical Issues* 15(1) (2005): pp. 386–405; Jeff McMahan, *Killing in War* (Oxford and New York: Oxford University Press, 2009a); 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, 2009b), pp. 385–394; Michael Otsuka, 'Killing the Innocent in Self-Defense', *Philosophy & Public Affairs* 23(1) (1994): pp. 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* (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): pp. 543–567.

² McMahan, *Killing in War*, p. 165.

³ Judith Thomson, *The Realm of Rights* (Cambridge MA: Harvard University Press, 1990), p. 229.

When it comes to drawing a principled line between an agent such as the conscientious driver and an agent such as the homeowner, I contend that we are facing a dilemma. On a first understanding of foreseeability, foreseeability is tied to fault, and on this understanding, *neither* the driver *nor* the homeowner could foresee the harm that they threaten. On a second understanding, foreseeability is likened to conceivability, and on this second understanding, *both* the driver *and* the homeowner could foresee the harms they threaten. I propose that we can escape this dilemma by adopting an understanding of foreseeable risk imposition that focuses on *duties of care*. In order to be morally permitted to drive, the conscientious driver had to abide by well-defined duties of care. No comparable duties applied to the homeowner. I argue that this renders the driver's actions foreseeably risk-imposing in a way that the homeowner's are not.

I then consider whether the Responsibility Account has an adequate rationale. Its proponents claim that it is *fairer*, and hence *distributively more just*, if an agent who is morally responsible for a threat of unjustified harm has to suffer that harm. They substantiate this claim by appealing to the luck egalitarian idea that it is distributively just to hold individuals to account for the consequences of their voluntary choices. Against this proffered justification, I argue that a commitment to holding individuals responsible cannot, by itself, ground foreseeable risk-imposers' defensive liability. More precisely, I show that it can help ground such liability only if independent moral reasons speak in favour of holding foreseeable risk-imposers accountable in this way. Moreover, an appeal to responsibility has justificatory force only if foreseeable risk-imposers are aware of the costs that they are asked to internalise, and only if their choices take place against a backdrop of substantive equality of opportunity. I conclude by suggesting that if the factors that render an agent liable to defensive harm remain constant across distributively just and unjust contexts alike, then what grounds such liability is plausibly not an agent's *responsibility* for threatening objectively unjustified harm, but her *culpability* for doing so.

The paper proceeds as follows. In section 2, I present the Responsibility Account. In section 3, I investigate how we might make sense of the notion of foreseeability that it relies on. In section 4, I discuss whether it is just to hold foreseeable risk-imposers to account for the harms that they threaten. Section 5 concludes.

II. THE MORAL RESPONSIBILITY ACCOUNT OF LIABILITY TO DEFENSIVE HARM

Consider the following cases:

The Conscientious Driver.⁴ A person who always keeps her car well-maintained and always drives carefully decides to drive to the cinema. On the way, her car veers out of control in the direction of a pedestrian, threatening the pedestrian with lethal harm. The pedestrian can avoid being run over by throwing a grenade towards the oncoming car. If this is the only way of defending himself, is it morally permissible for the pedestrian to kill the driver?

Day's End.⁵ A homeowner always comes home at 9 p.m. and flips the light switch in his hallway. He is about to do so this evening. The homeowner's flipping the switch causes a circuit to close. By virtue of an extraordinary series of coincidences, on this particular evening the circuit's closing will cause a lightning flash in a nearby house. If the homeowner flips the switch, the resulting lightning flash will kill his neighbour. If the neighbour can prevent the homeowner from flipping the switch only by killing him, is it morally permissible for the neighbour to kill the homeowner?

Proponents of the Responsibility Account argue that there is a crucial difference between the two cases that renders the driver, but not the homeowner, *liable to defensive harm*. The driver's liability then makes it permissible for the pedestrian to defend himself against the driver, whereas the homeowner's neighbour lacks a similar liability-based justification in favour of defending himself.⁶

If a person is *liable* to a particular harm, then they are not wronged if the harm is imposed on them. On the Responsibility Account, what grounds a person's liability to defensive harm is her *moral responsibility for a threat of harm that is objectively unjustified*.⁷ A threat of harm is *objectively unjustified* if it results from an action that lacks objective moral justification, so that it would have been impermissible to perform the action had the agent, at the time of acting, been aware of all the morally relevant facts, including all the morally relevant consequences of her alternatives for action. Importantly for our purposes, proponents of the Responsibility Account claim that a person can be morally responsible for a threat of unjustified harm without being blameworthy for it. On their understanding of moral responsibility, blameworthiness for a bad

⁴ McMahan, *Killing in War*, p. 165.

⁵ Thomson, *Realm of Rights*, p. 229.

⁶ Even if the homeowner is not liable to defensive harm, the neighbour may still be morally permitted to defend himself. For a defence of such a permission, see Jonathan Quong, 'Killing in Self-Defense', *Ethics* 119(3) (2009): pp. 507–537.

⁷ McMahan, *Killing in War*, p. 35.

outcome⁸ presupposes moral responsibility for that outcome, but when we act voluntarily and this leads to a bad outcome, we can sometimes be morally responsible for the outcome while being beyond reproach. More precisely, defenders of the Responsibility Account argue that an agent is morally responsible for an outcome of her action if, at the time of acting, the following conditions obtained:⁹

1. The agent was of sound mind;
2. The agent was in control of what she was doing;
3. It was foreseeable to the agent that her action might result in the kind of outcome that it actually resulted in.

I refer to the third condition above as the *Foreseeability Condition*. Defenders of the Responsibility Account think that the conscientious driver satisfies the Foreseeability Condition, whereas the unsuspecting homeowner does not.¹⁰ Here is what Jeff McMahan says about the conscientious driver:

‘She does not intend to harm anyone and cannot foresee that she will harm anyone, but she knows that driving is an activity that has a very tiny risk of causing great harm—so tiny that the activity, considered as a *type* of activity, is entirely permissible.’¹¹

The unsuspecting homeowner, by contrast, ‘does not choose to engage in an activity that has a foreseeable risk of causing serious harm’;¹² he is ‘nonculpably and invincibly ignorant’ of the fact that ‘he poses any kind of threat or risk of harm to anyone.’¹³

Neither the conscientious driver nor the unsuspecting homeowner is blameworthy for the harm that their actions threaten. But unlike the unsuspecting homeowner, the conscientious driver chose to engage in an activity that she knew, or should have known, imposed a risk of harm on innocent others. For defenders of the Responsibility Account, this fact about the driver’s epistemic situa-

⁸ In the context of this paper, I assume that a *threat of harm* qualifies as an *outcome of an action* even if the threat never materialises because it is successfully averted.

⁹ Otsuka, ‘The Moral-Responsibility Account’, p. 52.

¹⁰ *Ibid.*, p. 64.

¹¹ McMahan, *Killing in War*, p. 165, original emphasis.

¹² *Ibid.*, p. 166.

¹³ McMahan, ‘Self-Defense Against Morally Innocent Threats’, p. 394. In the passages just quoted, McMahan is not talking about the unsuspecting homeowner, but about an unsuspecting cell phone operator. I have replaced McMahan’s *Cell Phone Operator* (McMahan, *Killing in War*, p. 165) with Judith Thomson’s *Day’s End* because only the latter case is free from elements of intervening agency that might influence our judgments.

tion helps render the driver morally responsible for the harm that she is threatening, as it creates a morally relevant ‘association between the kind of act [that the driver voluntarily chose to perform] and the kind of threat [that she subsequently] poses.’¹⁴

According to defenders of the Responsibility Account, the conscientious driver’s moral responsibility for the threat that she poses makes her liable to defensive harm for reasons of ‘distributive fairness’.¹⁵ Their thought is that it is *fairer*, and hence distributively *less unjust*, if the driver has to suffer harm that is associated with her own voluntary actions than that the pedestrian has to suffer harm that is associated with someone else’s voluntary actions.¹⁶ As a morally responsible agent, the driver chose to engage in an activity that she ought to have known came with a risk of harm to innocent others. If the driver finds herself in a situation where this foreseeable risk of harm threatens to materialise, she cannot reasonably demand that the harmful consequences of her choice should be borne by someone other than herself. After all, the driver could have decided not to drive, thereby avoiding the risk of harm. It follows that when we are faced with a situation where either the driver or the pedestrian has to die, we do not wrong the driver if we ensure the pedestrian’s safety.¹⁷

When we turn our attention to *Day’s End*, no comparable justification is available for shifting the harm from the neighbour to the homeowner. The homeowner did not choose to engage in an activity that he knew, or should have known, imposed a risk of harm on innocent others. It is therefore no more just if the homeowner, as opposed to his neighbour, has to suffer harm.¹⁸

III. THE FORESEEABILITY DILEMMA

Intuitively speaking, defenders of the Responsibility Account seem right to claim that the conscientious driver was able to foresee that her actions might cause harm in a way that the unsuspecting homeowner was not. But the relevant difference between the two

¹⁴ McMahan, *Killing in War*, p. 168; see also Otsuka, ‘The Moral-Responsibility Account’, pp. 63–64.

¹⁵ Gordon-Solmon, ‘What Makes a Person Liable’, p. 546.

¹⁶ *Ibid.*; McMahan, *Killing in War*, pp. 206–207.

¹⁷ See also Jeff McMahan, ‘Who is Morally Liable to Be Killed in War?’, *Analysis* 71(3) (2011): pp. 544–559, p. 551.

¹⁸ Otsuka, ‘The Moral-Responsibility Account’, p. 64.

agents is difficult to articulate. When one tries to draw a principled distinction between the epistemic situation of the driver and that of the homeowner, it appears that either a threat of harm was relevantly foreseeable to *both* agents, or else it was relevantly foreseeable to *neither*. I refer to this difficulty of drawing a principled line between the two agents as the Foreseeability Dilemma.

A. *The No-Surprise View*

In one sense, neither the harm threatened by the conscientious driver nor the harm threatened by the unsuspecting homeowner was foreseeable. When they made their decisions to drive to the cinema and to switch on the light, both agents would rightly have been *surprised* to learn that their preferred alternative for action threatened significant harm. Otherwise they would both have been morally required to forgo trivial benefits to avoid threatening such harm.

Consider the case of driving a car. We are all aware that it is dangerous to drive if you have not received proper training, if you do not adjust your speed in bad weather, if you text while driving, or if you drive when sleepy or inebriated. Because of these dangers, we are morally required to exercise due care when driving. That is, we are not morally permitted to drive unless we have obtained proper training, are adjusting our driving to the prevailing weather conditions, and so on. If a driver violates driving-related duties of care in the absence of a morally weighty reason in favour of doing so, she acts negligently or recklessly, and her culpable behaviour makes her blameworthy.¹⁹ Negligent or reckless drivers are blameworthy partly because it comes as no surprise when their behaviour threatens harm. On this first and admittedly vague understanding of foreseeability, an agent thus *foreseeably* puts innocent others at risk if, and only if, *it comes as no surprise* if her actions threaten harm, and foreseeably putting others at risk is permissible only if there is a morally weighty reason in favour of doing so.

With respect to objectively unjustified harm that we risk threatening while pursuing small benefits, the just introduced No-Surprise View draws the line between foreseeable and unforeseeable threats of harm such that this line tends to track the distinction between

¹⁹ I use the following expressions interchangeably: being *blameworthy* for something, being *at fault* for something, and being *culpable* for something.

threats of harm that we are blameworthy versus not blameworthy for. The guiding thought behind the No-Surprise View is that *foreseeably putting innocent others at risk is pro tanto wrong*, and that ‘foreseeably’ needs to be interpreted in a way that is consistent with this verdict.

Something like the No-Surprise View of foreseeability plays an important role in negligent liability.²⁰ For one thing, it is already a well-established practice to apply a test of *reasonable foreseeability* to determine whether a defendant’s breach of a duty of care should count as the *proximate cause*²¹ of some damage that a plaintiff incurred. That is, only if it is deemed reasonably foreseeable—or non-surprising—that the defendant’s breach of their duty should have ‘cause[d] the type of hazard to the type of person’²² that it actually did cause may the breach be seen as a proximate cause of the damage, and the defendant thus be held liable for the damage (see also fn. 39 below). In addition, it is increasingly recognised that the idea of reasonable foreseeability also plays an important role in determining whether an agent was, in the first place, breaching any duties of care.²³

Theorists of tort law frequently criticise the notion of foreseeability for its vagueness. The recurrent worry is that foreseeability is so ‘notoriously open-ended [that it] can be used to explain any decision, even decisions directly opposed to each other.’²⁴ Those who defend the use of the notion admit that it is hard to pin down, but argue that it nevertheless plays an ineliminable role in our judgments about who may justly be held to account for which damages.²⁵

We can leave open whether foreseeability in the No-Surprise sense is sufficiently determinate to play a productive role in tort law. For our purposes, it matters that it is sufficiently determinate to

²⁰ See, e.g., Benjamin Zipursky, ‘Foreseeability in Breach, Duty, and Proximate Cause’, *Wake Forest Law Review* 44 (2009): pp. 1247–1275.

²¹ David Owen follows Henry W. Edgerton in defining a proximate cause as a ‘justly-attachable cause.’ See David G. Owen, ‘Figuring Foreseeability’, *Wake Forest Law Review* 44 (2009): pp. 1277–1307, p. 1300, fn. 119.

²² *Ibid.*, p. 1291.

²³ Zipursky, ‘Foreseeability’, esp. pp. 1255–1257.

²⁴ Patrick Kelley, ‘Restating Duty, Breach, and Proximate Cause in Negligence Law: Descriptive Theory and the Rule of Law’, *Vanderbilt Law Review* 1039 (2001): pp. 1039–1070, p. 1046. For a summary of critical voices, see Owen, ‘Figuring Foreseeability’, pp. 1277–1280.

²⁵ See, e.g., Owen, ‘Figuring Foreseeability’, p. 1280.

conclude that *neither* the harm threatened by the conscientious driver *nor* the harm threatened by the unsuspecting homeowner falls under its scope. Both the driver and the homeowner would rightly have been surprised to learn that their actions threaten harm, and neither agent was at fault for pursuing trivial benefits in a way that turned out to threaten disaster. Yet there nevertheless remains a sense in which the possibility of a harmful outcome *was* epistemically accessible *at least to the driver* at the time of acting. We all know that even the best maintained cars can still malfunction, and that through a series of unfortunate events, this can create a threat of harm to innocent others. The next section discusses an interpretation of foreseeability that captures this thought.

B. *The Conceivability View*

If we use the term ‘foreseeable’ as roughly interchangeable with terms such as ‘possible’ or ‘conceivable,’ then it is clearly foreseeable that even a conscientious driver might cause harm. This is so because we can always come up with stories to explain why even careful drivers *might* cause accidents. From an *ex ante* perspective, it therefore seems reasonable consistently to assign a non-zero probability to harm occurring as a result of anyone’s decision to drive, as well as to harm occurring as a result of people’s actions more generally.

David Owen expresses the thought behind this *Conceivability View* of foreseeability, and, by extension, the Foreseeability Condition of moral responsibility, as follows:

‘If choice is so important [to holding people morally responsible for an outcome that they caused], then why not ground responsibility simply on the choice to *act*—knowing, as we all do, that unexpected and even unforeseeable results, if uncommon, not infrequently do occur. The argument here might be that *unforeseeable* consequences, those types of consequences that people comprehend as lying outside the realm of normal expectation but that sometimes do occur, are then actually *themselves* foreseeable [...].’²⁶

In their seminal work *Causation in the Law*, H.L.A. Hart and Tony Honoré sketch an account of moral responsibility that they refer to as *outcome responsibility*.²⁷ The core idea of their account—later fleshed out in more detail by Honoré²⁸—resembles the thoughts

²⁶ Owen, ‘Figuring Foreseeability’, p. 1283.

²⁷ H.L.A. Hart and Tony Honoré, *Causation in the Law* (Oxford and New York: Oxford University Press, 1985).

²⁸ See Tony Honoré, ‘Responsibility and Luck. The Moral Basis of Strict Liability’ in *Responsibility and Fault* (London: Bloomsbury Publishing, 1999), pp. 14–40.

voiced by Owen in the quote above. Outcome responsibility is a form of strict responsibility that is independent of fault; an agent acquires outcome responsibility for a harm simply by bringing it about through her voluntary actions. Hart and Honoré argue that, for better or worse, our identities are shaped by the outcomes of our voluntary actions, and that this renders it appropriate that others hold us accountable for these outcomes.²⁹

While Hart and Honoré's account of moral responsibility may have its merits, it sits ill with the Responsibility Account for at least two reasons. First, the Responsibility Account stipulates that a person's moral responsibility for a threat of objectively unjustified harm suffices to make that person liable to potentially lethal defensive harm. Yet Honoré is clear that an agent's outcome responsibility for another's injury is insufficient, by itself, to ground liability to pay compensation. In Honoré's words, when conduct that is not 'socially undesirable' results in harm, then '[a]n [agent's] apology [to their victim] will often be enough.'³⁰ If this is correct, then outcome responsibility is too weak a notion to ground defensive liability.

Second, if we construe foreseeability as conceivability—as we are pushed to do if we accept Hart and Honoré's account of moral responsibility, yet still want to hang on to a Foreseeability Condition—then *both* the conscientious driver *and* the unsuspecting homeowner are responsible for the harm that they threaten. While it is *extraordinary* to threaten one's neighbour merely by flipping a light switch, such a freak event nevertheless remains possible.

It follows that there is a dilemma for the Responsibility Account. On a first intuitive understanding of foreseeability, foreseeable harm is unsurprising harm that an agent has a *pro tanto* reason to avoid, and neither the driver nor the homeowner could foresee the harm they threaten. On a second intuitive understanding, foreseeability is likened to conceivability, and on this understanding, harm was foreseeable to both agents. Moreover, it is not obvious that there is a further understanding of foreseeability that allows us to draw a principled distinction between the two agents.

²⁹ Hart and Honoré, *Causation in the Law*, p. lxxx. See also John Gardner, 'Obligations and Outcomes in the Law of Torts' in P. Cane and J. Gardner (eds.), *Relating to Responsibility. Essays for Tony Honoré on his Eightieth Birthday* (Oxford: Hart Publishing, 2001), pp. 111–144, pp. 133–134.

³⁰ Tony Honoré, 'The Morality of Tort Law—Questions and Answers' in D. Owen (ed.), *Philosophical Foundations of Tort Law* (Oxford: Clarendon Press, 1997), pp. 73–95, p. 82.

C. *Escaping Between the Horns of the Dilemma*

We take a first step towards solving the Foreseeability Dilemma once we realise that the difference between the conscientious driver and the unsuspecting homeowner might be closer to a difference in degree rather than in kind. While the epistemic accessibility of outcomes that come as no surprise is clearly distinguishable from the accessibility of merely conceivable outcomes, it nevertheless seems true that the No-Surprise View and the Conceivability View merely delimit two different segments on a *continuum of epistemic accessibility of outcomes*. This continuum ranges from outcomes that an agent, at the time of acting, reasonably perceives to be *certain* to outcomes that she reasonably perceives to be *impossible*.

1. *The Quantitative Threshold Approach*

If we accept the idea of a continuum of epistemic accessibility, it looks like settling for a *quantitative threshold* might provide a way out of the Foreseeability Dilemma. The thought behind such an approach is that the morally relevant difference between the driver and the homeowner is that the former imposes a *larger risk* on others than the latter. Quantitative threshold solutions thus stipulate that there is a cut-off point at which the risks of harm that we impose on others become so vanishingly small that the activity that we are engaged in no longer counts as foreseeably risk-imposing. This cut-off point is then thought to get the homeowner, but not the driver, off the hook.

For two reasons I believe that quantitative approaches are not promising. The first reason is epistemic. If there is a cut-off point below which our activities cease to be foreseeably risk-imposing, then we should at least usually be able to establish whether we would lie above or below this cut-off point if we decided to engage in some activity. To be sure, if the concept of foreseeability marks out a segment on a continuum, then it is not surprising that for any specification of the concept, there will exist borderline cases where we are unsure whether an outcome should count as foreseeable. But it is a problem for the Responsibility Account if agents *as a matter of course* have trouble establishing whether some activity would count as foreseeably risk-imposing. As defenders of the Responsibility Account make clear, if an agent engages in a foreseeably risk-imposing activity, this means that there is a morally relevant 'association be-

tween the kind of act [the agent chooses to perform] and the kind of threat [she subsequently] poses.³¹ This suggests that a foreseeably risk-imposing activity is an activity that an agent can usually discern, and hence choose to perform, as such. Yet if we take the idea of a quantitative cut-off point seriously, it is difficult to know whether an agent falls above or below it *whenever she exercises a potentially dangerous activity while exercising due care*.

For illustrative purposes, suppose that an activity is foreseeably risk-imposing if there is a higher than one in a million probability that pursuing the activity for one hour will lead to substantial harm to an innocent other. If I have a gentle dog and I am an experienced dog person, does the probability of harm that I impose on others by walking my dog fall above or below this threshold? Even if the available data suggests that ‘dog walking’ is a foreseeably risk-imposing activity, the fact that I am *experienced* and my dog is *gentle* might mean that *my* dog-walking imposes only morally insignificant risks. There are further factors that seem relevant to determining what risks I impose on others, such as where I walk my dog, whether I keep him on a leash at all times, and what breed my dog is. Even if I know where the relevant cut-off point for foreseeability lies, it is extremely unlikely that I will have access to sufficiently fine-grained data to estimate accurately, and in good faith, whether I exceed it when walking my dog.

To make matters worse—and this is the second reason against adopting a quantitative threshold approach to solving the Foreseeability Dilemma—the difficulties involved in quantifying risks are not merely epistemic. If we assume that the size of the risk that my dog-walking imposes is *belief-independent*, in the sense that different reasoners, when presented with the same facts, should all reach the same conclusion about the level of risk imposed, then there will not usually be a uniquely right conclusion that they all should reach.

The most useful belief-independent interpretation of what it means to impose a probability of harm is offered by *finite frequentism*. Roughly speaking, finite frequentism defines the probability of a particular harm as the relative frequency with which the harm occurs in a relevant reference class. If we take all the experienced dog walkers with gentle dogs who walked their dogs in 2018—assuming

³¹ McMahan, *Killing in War*, p. 168.

we can render the terms ‘experienced dog walker’ and ‘gentle dog’ sufficiently determinate—and we divide the number of harmful incidents by the total number of dog walking hours, this yields the frequentist probability of harm *for this particular reference class*. The problem with this intuitively appealing procedure is that there is no uniquely fitting reference class for any particular instance where an agent is contemplating whether to engage in some activity in some particular manner.³² The risk of harm that I impose on others qua ‘experienced dog walker who walked their gentle dog in 2018’ may be quite different from the risk that I impose on others qua ‘experienced dog walker who walked their gentle dog between 2008–2018’. This so-called *reference class problem* means that frequently there may be no fact of the matter about the size of the risk that an agent imposes if she chooses to engage in some activity in a particular manner.³³ A quantitative threshold approach thus does not help us pin down a conspicuous ‘association between the kind of act [a foreseeably risk-imposing agent chooses to perform] and the kind of threat [she subsequently] poses.’³⁴

2. *The Duties of Care Approach*

Luckily, I believe that we can do better. I propose that we should think of activities as *foreseeably risk-imposing*, or as being of the *risk-imposing type*,³⁵ just in case performing them is morally permissible in the evidence-relative sense³⁶ only if agents abide by relatively well-defined *duties of care*. On this proposal, driving a car or walking a dog are *always* foreseeably risk-imposing activities, simply because we are

³² See, e.g., Alan Hájek, ‘The Reference Class Problem Is Your Problem Too’, *Synthese* 156 (2007): pp. 563–585.

³³ Someone attracted to the quantitative approach might jump to its defence here. While the reference class problem frequently makes it impossible to precisely quantify the size of the risk that an agent imposes, it may in many such cases still be possible to *delimit a set of plausible reference classes*, and to calculate what risks an agent imposes as a member of these different reference classes. It might then make sense to refer to an activity as foreseeably risk-imposing if, for a sufficient proportion of relevant reference classes, the risk that an agent imposes falls above the relevant cut-off point. I grant that such strategies help alleviate the metaphysical issues just introduced, but only at the cost of exacerbating the epistemic worries raised earlier.

³⁴ McMahan, *Killing in War*, p. 168.

³⁵ McMahan introduces this second term as an alternative to *foreseeably risk-imposing* in ‘The Basis of Moral Liability’, p. 397.

³⁶ Derek Parfit refers to an act as morally permissible in the *evidence-relative sense* ‘just when this act would be [morally permissible] in the ordinary sense if we believed what the available evidence gives us decisive reason to believe, and these beliefs were true.’ See Derek Parfit, *On What Matters, Volume 1* (Oxford and New York: Oxford University Press, 2011), p. 151.

morally required to guard against certain dangers when performing them.

The general thought behind this proposal is as follows. Over time, we come to understand that certain activities are associated with dangers. This understanding might be based simply on the *observation of correlations*—we might for example observe that individuals who like to experiment with chemicals frequently start uncontrolled fires in their backyards.³⁷ Alternatively, our understanding might run deeper: we might appreciate how certain *causal processes* that are associated with the activity can lead to harmful outcomes unless carefully guarded against. To continue with the aforementioned example, we might understand what chemical reactions lead to combustion, as well as the conditions under which fires will spread. When we come to understand that an activity is associated with certain dangers, it remains morally permissible to perform the activity in the pursuit of our daily lives if, and only if, two conditions obtain. First, the activity serves a morally valuable end, in the sense that being able to engage in it contributes to our ability to lead flourishing lives.³⁸ Second, we find ways of managing the dangers associated with the activity by imposing certain standards of due care. I contend that it is this subset of activities—the activities that we may permissibly perform only if we abide by certain duties of care—that should count as *foreseeably risk-imposing*.

A strength of this proposed solution is that it establishes a clear connection between the kind of act that a foreseeably risk-imposing agent chooses to perform and the harm she may subsequently be held morally responsible for. If an agent performs an activity that comes with duties of care, then even if the agent goes to great lengths to ensure that she does not put innocent others at risk, she

³⁷ The two problems raised for the quantitative threshold approach (see section 3.3.1) do not stand in the way of this kind of insight. The *epistemic problem* does not arise because we rightly disregard the *manner in which an activity is performed* when we observe correlations. When we come to understand that an activity is dangerous based on the harmful outcomes that it is correlated with, we simply observe the people who engage in the activity, and deem the activity dangerous if, and only if, engaging in it is associated with a harmful outcome sufficiently frequently. The *metaphysical problem* does not arise because we are free to *define*—sensibly by following common language use—what counts as *performing the activity*. Simply stipulating what counts as performing the activity is admissible as long as we refrain from claiming that every agent who performs the activity thereby imposes a particular risk of harm. To put the same point differently, when we call an activity dangerous based on observed correlations, we make a statement about the *average* person performing the activity, while allowing that there might be large variations between individuals. I thank an anonymous reviewer for pressing me on this issue.

³⁸ See John Oberdiek, *Imposing Risk* (Oxford and New York: Oxford University Press, 2017), pp. 131–154.

still has reason to assume that she thereby only minimises, but does not fully eliminate, the dangers associated with the activity. Even if there is no fact of the matter about the size of the risk that an agent imposes if she engages in a potentially dangerous activity while abiding by relevant duties of care, the fact that the agent is morally required to exercise due care gives us reason to believe that the activity comes with risks that an innocuous activity lacks. This provides the agent with a—possibly only very weak—moral reason in favour of avoiding the activity. Moreover, it is implied by the idea of duties of care that an agent *should be aware* that such duties apply to her, as the agent is *at fault* if she violates her duties by behaving negligently or recklessly. If a foreseeably risk-imposing activity is an activity that an agent may permissibly perform only if she exercises due care, an agent should thus be able to recognise foreseeably risk-imposing activities as such.

But the proposed solution also calls for clarification and refinement. First, note that there are certain dangerous activities the dangers of which we are not yet aware, such that there are no duties of care that a reasonable agent would recognise as being associated with these activities. For example, suppose that inhaling someone else's bad breath can cause lung cancer, though we currently lack medical evidence for this. If this were the case, then talking to people when you have bad breath would be an example for an activity that is dangerous in this unidentified manner. I contend that activities of this type should *not* count as foreseeably risk-imposing. An activity is *foreseeably* risk-imposing only if it is reasonably epistemically accessible to an agent at the time of acting that engaging in the activity will expose others to dangers that the agent can manage, but not completely eliminate. Only if this condition is satisfied is the activity intelligibly avoidable for being of the risk-imposing type, so that it might be fairer to shift the potential costs of the activity on the agent who decided to undertake the activity as opposed to letting them fall where they may.

Second, suppose that there is an activity that comes with certain duties of care because of its associated dangers, and an agent decides to perform the activity while exercising due care. Further suppose that this agent's conduct now threatens unjustified harm, but the reason why it does so is not related to the dangers associated with

the activity. For example, suppose that a recovering heart patient is sitting on his porch when you drive by in your car (at a steady and moderate speed, and on a road where driving is permissible). As it happens, the noise of your engine startles the patient, and he has a heart attack. Neither the patient nor you had any reason to expect this outcome. It seems to me that in such a case, we cannot hold you morally responsible for the harm that you caused. The harm that materialises is no more associated with your driving than the harm that threatens to materialise in *Day's End* is associated with the homeowner's decision to flip the switch; both harms are freak events that an agent has no reason to attend to when deciding whether to engage in the relevant activities. If this is correct, then we should say that according to the Responsibility Account, an unexcused agent is *morally responsible* for a threat of harm that her *foreseeably risk-imposing* activity causes if, and only if, the threat of harm is due to a danger that she was morally required to manage by exercising due care.³⁹

With these two refinements, we get the following definitions:

An activity is of the **risk-imposing type** or **foreseeably risk-imposing** if, and only if, performing the activity is morally permissible in the evidence-relative sense only if an agent abides by certain duties of care.

An otherwise unexcused agent is **morally responsible for a threat of objectively unjustified harm** that she poses by engaging in a foreseeably risk-imposing activity if, and only if, the threat of harm is due to a danger that is associated with the activity. A danger is associated with an activity if some of the duties of care that accompany the activity are in place to guard against this danger.

As I see it, these definitions help flesh out the Responsibility Account in a plausible manner. They clarify that there is a principled and potentially morally relevant difference between the conscientious driver and the unsuspecting homeowner. In what follows, I investigate whether the difference between the two agents is morally significant in the way that defenders of the Responsibility Account claim it is.

³⁹ In tort law, a similar distinction is made, in that defendants are held liable for harms only if these harms are considered to be *within the risk* of a particular conduct (strict liability) or breach of duty (negligent liability). See section 3.1 above.

IV. THE MORAL SIGNIFICANCE OF ENGAGING IN A FORESEEABLY RISK-IMPOSING ACTIVITY

A. *Voluntary Choice and Control*

Defenders of the Responsibility Account argue that if a fixed amount of harm has to be suffered either by a foreseeable risk-imposer or her potential victim, then it is *fairer* if the harm is suffered by the risk-imposer. Their thought is that a foreseeable risk-imposer, unlike her potential victim, had at least some morally relevant *control* over the threat's existence.⁴⁰ As McMahan puts it with respect to the conscientious driver,

[T]he driver ha[s] made herself liable by *choosing* to set a couple of tons of steel in motion as a means of pursuing her ends, knowing that this would involve a very small risk of killing an innocent person. Again, a *voluntary choice* with a foreseeable risk [...] is the basis of [defensive] liability.⁴¹

This quote suggests that the driver had control over the threat's existence because she made a *voluntary choice* to engage in a foreseeably risk-imposing activity. But what makes her choice voluntary? First—and this is something McMahan makes explicit—it must have been *epistemically accessible* to the agent that she was choosing to engage in a foreseeably risk-imposing activity. Second—and this is left implicit in the passage just quoted—the agent must have had access to *reasonable alternatives for action* that would not have been similarly risk-imposing. Only if she had access to such alternatives was it reasonably avoidable for her to engage in an activity of the risk-imposing type.⁴² It is difficult to determine what makes an alternative reasonable to choose. For the problem at hand, I suggest the following: if it is morally permissible to engage in some activity, then available alternatives are *reasonable to choose* if, and only if, they are not significantly inferior in terms of the opportunities for welfare they offer to the agent.⁴³ By contrast, if engaging in an activity is morally impermissible, then permissible

⁴⁰ See, e.g., Otsuka, 'The Moral-Responsibility Account', p. 63; McMahan, 'Self-Defense Against Innocent Threats', p. 392.

⁴¹ McMahan, *ibid.*; emphases added.

⁴² This second condition is mentioned e.g. in Otsuka, 'Killing the Innocent', p. 91, and is discussed, albeit inconclusively, in Otsuka, 'The Moral-Responsibility Account', pp. 63–65. See also Gordon-Solmon, 'What Makes a Person Liable', pp. 558–559.

⁴³ More precisely, I assume that alternatives for action are reasonable to choose whenever they are not significantly inferior when evaluated in terms of the *right currency of distributive justice*, which for ease of exposition I assume to be *opportunities for welfare*. None of my arguments depend on this assumption.

alternatives are always reasonable, even if they offer clearly inferior opportunities for welfare to the agent.

On this specification of what it means for a choice to be voluntary, it seems plausible that the driver's voluntary choice to drive might ground her defensive liability, as there is an intelligible sense in which the driver controls whether a threat of objectively unjustified harm might come into existence. Proponents of the Responsibility Account appeal to luck egalitarian ideas to substantiate their claim that it is just to hold agents to account for the consequences of their voluntary choices.⁴⁴ In what follows, I turn to luck egalitarianism and argue that, on closer inspection, it does not generally support the claim that it is fairer if foreseeable risk-imposers, as opposed to their potential victims, have to absorb the harms that they threaten.

B. Luck Egalitarianism

The basic tenet of luck egalitarian theories of distributive justice is that 'it is unjust if some people are worse off than others through [sheer] bad luck.'⁴⁵ If a person is worse off than others due to her voluntary choices, then she is worse off than others for reasons other than her bad luck. Seemingly very much in line with the justificatory rationale behind the Responsibility Account, luck egalitarians are committed to the idea that a key reason to favour equality—namely the unfairness of differential luck—is absent when people's differential fates are the result of factors within their control.

The starting point of all luck egalitarian theories of distributive justice is one of substantive, or fair, equality of opportunity.⁴⁶ For luck egalitarians, it is a requirement of justice that everyone's choice set—the options that people choose from—is equally valuable in terms of the opportunities for welfare that it provides (cf. fn. 43). Against such a background of equal opportunity, luck egalitarians are then committed to the idea that it is just to hold people to account for their voluntary choices, thus respecting them as responsible agents.⁴⁷ To put

⁴⁴ See Gordon-Solmon, 'What Makes a Person Liable'; Otsuka, 'The Moral-Responsibility Account', pp. 63–64.

⁴⁵ Kasper Lippert-Rasmussen, *Luck Egalitarianism* (London: Bloomsbury Publishing, 2016), p. 1.

⁴⁶ See, e.g., Richard Arneson, 'Equality and Equal Opportunity for Welfare', *Philosophical Studies*, 56 (1989): pp. 77–93.

⁴⁷ See, e.g., G.A. Cohen, *On the Currency of Egalitarian Justice* (Princeton: Princeton University Press, 2011), p. 4.

the same point differently, if some people end up worse off than others through their voluntary choices from equally valuable choice sets, then there is no demand of justice to eliminate the resulting inequalities. To the contrary, it would be unjust if those who made more prudent choices were made worse off in order to improve the fate of those who were imprudent. The situation is different if someone's opportunities for welfare are reduced through sheer bad luck. In this case, it is unjust that the unlucky person is worse off than others.

Because inequalities are considered unjust if they are due to bad luck, it might seem that on a luck egalitarian picture, it is just as unjust if the conscientious driver as opposed to the pedestrian has to suffer harm. After all, the driver is unlucky: most conscientious driving does not result in a threat of harm. While this worry is important, luck egalitarians have a response to it. The response draws on Ronald Dworkin's distinction between *brute luck* and *option luck*. According to Dworkin, it is unjust if some people are worse off than others through bad brute luck, but it need not be unjust if some people are worse off than others as a result of their bad option luck. As Dworkin defines it, *option luck* relates to foreseeable outcomes that are the results of risky actions that an agent could reasonably have avoided, whereas *brute luck* relates to outcomes an agent could not have similarly avoided:

'Option luck is a matter of how deliberate and calculated gambles turn out—whether someone gains or loses through accepting an isolated risk he or she should have anticipated and might have declined. Brute luck is a matter of how risks fall out that are not in that sense deliberate gambles. If I buy a stock on the exchange that rises, then my option luck is good. If I am hit by a falling meteorite whose course could not have been predicted, then my bad luck is brute [...].'⁴⁸

The main reason for distinguishing between brute luck and option luck is as follows. To respect people as responsible agents, a society needs to provide them with access to meaningful options. This includes providing them with access to *gambles*, i.e. alternatives for action where it is epistemically accessible to the agent that things could turn out better or worse, and where the agent is able to assign at least rough probability estimates to the likelihood with which different outcomes might occur. When an agent has access to a gamble, and the agent chooses to engage in it even though she could reasonably have declined it for a safe alternative, then it is not unjust

⁴⁸ Ronald Dworkin, 'What is Equality? Part 2: Equality of Resources', *Philosophy & Public Affairs* 10 (1981): pp. 283–345, p. 293.

if the gambler ends up worse off or better off than other gamblers, just as it is not unjust if the gambler ends up worse off or better off than those who chose a safe alternative instead. If the inequalities that result from gambles were seen as unjust and in need of elimination, then the redistribution that would have to take place in the name of distributive justice would erode the option to gamble, thus restricting the choices people have access to.

Applied to the notion of a foreseeably risk-imposing activity, the thought here is that engaging in a foreseeably risk-imposing activity means choosing to gamble. When engaging in such an activity threatens harm, the agent suffers bad option luck, and it is not unjust to fail to correct this. Importantly, though, this is a valid way of looking at things from a luck egalitarian perspective only if a number of demanding background conditions obtain.

C. Equality of Opportunity and Reasonable Avoidability

From McMahan's description of the case, we know very little about the conscientious driver and the pedestrian whose life she threatens. But suppose the following details are correct: the driver and the pedestrian both live in a small town. The driver has a disability that makes it painful for her to walk. Because of its associated risks, she dislikes driving. Hence she always thinks twice about leaving her house, and goes to great lengths to drive safely. The pedestrian, by contrast, has no disability. He loves walking and cycling, and does not even have a driver's licence. Even in the absence of the option to drive, his choice set is more valuable to him in terms of opportunities for welfare than the driver's choice set is to the driver.

Fleshing out the situation in this way clarifies that it is generally⁴⁹ fairer if harm is borne by the person who created an unjustified threat of harm by engaging in a foreseeably risk-imposing activity only if everyone enjoys equal access to opportunities for welfare. In the absence of such equal opportunities, it is always possible that some people will be able to reach a certain level of welfare without engaging in any foreseeably risk-imposing activities, whereas others

⁴⁹ The qualifier 'generally' is needed because it is always possible that the pedestrian's choice set is inferior to the driver's. When this is the case, it will be fairer to shift harm from the pedestrian to the driver. But this will be fairer not primarily because the driver is responsible for the harm that she threatens, but because it equalises opportunities for welfare between the driver and the pedestrian if the harm is apportioned to the driver.

will be able to reach a comparable level of welfare only if they choose to engage in activities of the risk-imposing type.

Describing the driver as disadvantaged also helps clarify the importance of reasonable avoidability. Suppose that if the driver engages in driving, then on expectation she is as well off as the pedestrian, who does not drive. But if she doesn't drive, then she is much worse off. In this case, driving is not reasonably avoidable for her, because she does not have access to safe alternatives to driving that present her with similarly valuable opportunities for welfare. In the absence of such alternatives, it does not seem just to burden the driver with the costs of a gamble that she would gladly have passed up in favour of a safe alternative.

In sum, we are entitled to conclude that engaging in a foreseeably risk-imposing activity can be described as *taking a gamble that agents may justly be held responsible for* only if the background conditions of equal opportunity and reasonable avoidability obtain.

D. Specified Gambles and Shared Expectations

Imagine a *homogeneous society* where all agents are the same in terms of their abilities, their preferences, and their choice sets. In this society, substantive equality of opportunity is realised, as everyone has access to the same opportunities for welfare. Suppose that in this society, driving is a reasonably avoidable gamble. Compared to the safe alternative of taking public transport, it is more convenient, but it comes with risks of harm that drivers are asked to internalise. In terms of expected welfare, driving is slightly more attractive than its safe alternative, but the difference is small. Now suppose that in this society, everyone chooses to drive to a similar extent. Most people drive without ever creating a threat of unjustified harm. But now suppose that McMahan's conscientious driver is a member of this society, and that one fine evening, her decision to drive threatens a pedestrian. Is it correct to say that it is fairer if the driver, as opposed to her potential victim, should die as a result of the driver's choice?

Here is what the pedestrian might say, inspired by the justification behind the Responsibility Account: 'Yes, it is less unjust if the driver dies. She's made a voluntary choice to drive when she could have taken the subway instead. When she decided to drive, she knew that driving is a risk-imposing activity. Because she chose to engage in an avoidable gamble, it is not unjust if we hold her to account for the outcome of this

gamble. At the very least, it is less unjust if she is worse off as an outcome of a gamble she took than if I am worse off as an outcome of a gamble she took.' In the context of a *homogeneous* society, these words ring hollow. True, the driver chose to drive. But the pedestrian would have made a similar decision in her place. By assumption, both the driver and the pedestrian regularly choose to drive, and for understandable reasons: the expected benefits of doing so outweigh the expected costs. It is true that there is a threat of harm only because the driver decided to drive to the cinema on this particular evening. In this sense, it can be argued that a natural lottery has taken place which has singled out the driver: it just so happens that *her* decision to drive has resulted in a threat of objectively unjustified harm, and we may as well hold her to account for this. But if this is all that we can say, then a similar argument can be made with respect to the unsuspecting homeowner: it just so happens that his decision to flip the light switch has resulted in a threat of objectively unjustified harm, and we may as well attach the harm to the person whose choice has created the threat. In other words, in the homogeneous society that I have just sketched, the difference between the driver and the pedestrian is reduced to the fact that the former's actions threaten objectively unjustified harm, whereas the latter's do not. With respect to their voluntary choices, the driver and the pedestrian are indistinguishable.⁵⁰

To see this crucial point more clearly, consider Dworkin's example of investing on the stock market. Suppose that two investors both conduct a similar amount of research and then decide to invest part of their wealth in stocks instead of keeping it in a savings account. Further suppose that one investor makes a profit while the other has to take a loss. In such a case, there is no sense in which it is *more just* if we leave the situation as it is instead of reversing it. If there is a reason why we should leave the situation as it is, it is tied to considerations other than a commitment to holding people responsible for their choices (e.g. efficiency), as both investors' choices were relevantly the same.

Of course, the driver and the pedestrian *needn't* be the same in terms of their voluntary choices. Once we move away from the assumption of a perfectly homogeneous society, there are certain conditions under which it is less unjust if the driver has to suffer

⁵⁰ In tort law, George Fletcher argues for a regime of negligence (as opposed to strict) liability where risk imposition is symmetrical for reasons that mirror the reasons presented in this paragraph. See George Fletcher, 'Fairness and Utility in Tort Theory', *Harvard Law Review* 85 (1972): pp. 537–573.

lethal harm. But surprisingly it might also be the case that it is less unjust if the pedestrian has to suffer harm. To see what I mean, suppose that in a society with heterogenous agents, we can design two different sets of rules. Under the first set of rules, we ask drivers to absorb all driving-related costs. Under the second set of rules, we determine that driving-related costs which are the result of conscientious driving should fall where they may. Further suppose that both sets of rules, thanks to further differences in their specifications, are compatible with substantive equality of opportunity, and would present agents only with reasonably avoidable gambles. Finally, suppose that no matter which set of rules we implement, some people will decide to drive frequently, others will decide to drive rarely, and again others will rely exclusively on public transport. Which set of rules should we implement *as more just*?

The crucial point is that luck egalitarianism is indifferent: it regards the two sets of rules as *equally just*. It is true that depending on which set of rules we implement, the members of our society will have access to different options, as the rules that we implement specify the gambles that individuals have access to. If we implement the first set of rules, then choosing to be a driver will be a slightly less attractive gamble; if we implement the second set of rules, then this renders walking slightly less attractive. But luck egalitarian principles do not help us determine which options it is morally most valuable to have access to. As Serena Olsaretti puts it, '[...] even if we accept the importance of holding individuals responsible for their choices, it is an open question what consequences those choices should have, and [our view on what consequences individuals' choices should have] presupposes substantive moral convictions that are independent of our commitment to responsibility itself.'⁵¹

It follows that we cannot appeal to luck egalitarian ideas to settle the question whether agents should be held accountable for whatever threats of objectively unjustified harm that they pose through

⁵¹ Serena Olsaretti, 'Responsibility and the Consequences of Choice', *Proceedings of the Aristotelian Society, New Series*, 109 (2009): pp. 165–188, p. 167. Zofia Stemplowska argues similarly that no luck egalitarian theory is complete unless its commitment to responsibility is complemented with principles that determine what options agents should have access to. See Zofia Stemplowska, 'Making Justice Sensitive to Responsibility', *Political Studies* 57 (2009): pp. 237–259, and Zofia Stemplowska, 'Rescuing Luck Egalitarianism', *Journal of Social Philosophy*, 44(4) (2013): pp. 402–419. Others who agree include Shlomi Segall, *Health, Luck, and Justice* (Princeton: Princeton University Press, 2010), ch. 4 and Alex Voorhoeve, 'May a Government Mandate More Comprehensive Health Insurance than Citizens Want for Themselves?', *Oxford Studies in Political Philosophy*, 4 (2018): pp. 167–191, p. 178.

engaging in foreseeably risk-imposing activities. The question whether agents should be held responsible in this way is logically independent of luck egalitarian justice, in the sense that however we answer it, this answer will be compatible with the luck egalitarian commitment to holding people responsible for their choices. Whether foreseeable risk-imposers should have to internalise the costs of their behaviour might depend, for example, on what would promote the general welfare.⁵²

Importantly, luck egalitarianism is *not* indifferent on whether we should hold people responsible for the choices *that they were presented with*. From a luck egalitarian perspective, it is important that agents are aware of the choices in their choice set—that they know their alternatives for action, and that it is clear to them what consequences are attached to which alternatives. In fact, it is precisely the point of the luck egalitarian commitment to holding people responsible that we all have to live with the consequences of our choices *in line with what we reasonably expected these consequences to be*. It follows that whichever set of rules we decide to implement, it is important that agents are made aware of the rules that apply to them. If agents in a society characterised by equal opportunity and reasonable avoidability of gambles know that pedestrians are asked to absorb certain driving-related costs, then it is precisely this fact that renders it just that pedestrians, as opposed to drivers, have to absorb these costs.

At this point, it might be objected that something has gone wrong. Even if it is possible to attach driving-related harms to pedestrians, thus turning walking into a less attractive option than it would otherwise be, there seems to be something unjust about burdening pedestrians with such harm. After all, it seems that if pedestrians have to absorb driving-related harms, then drivers are permitted to inflict non-consensual harm on pedestrians. This seems highly counterintuitive. But note that drivers are *not* permitted to

⁵² From a perspective that is concerned with promoting the general welfare, economists argue that asking agents to internalise the costs of their self-interested activities helps fix the extent to which these activities are pursued at its social optimum. This provides a welfarist consideration in favour of implementing the set of rules that the Responsibility Account advocates, at least with respect to the foreseeably risk-imposing activities that agents undertake mostly for their own sake. Having said that, it is easily imaginable that it would detract from the general welfare if people started arming themselves so as to be able to defend themselves against perceived threats posed by others' self-interested risk-imposing activities (see Seth Lazar, 'Responsibility, Risk, and Killing in Self-Defense', *Ethics* 119(4) (2009): pp. 699–728, pp. 718–722). These and other welfare effects would have to be estimated and balanced against each other.

inflict non-consensual harm. What they *are* permitted to do is *engage in a foreseeably risk-imposing activity while exercising due care*. Predictably, this will sometimes result in a threat of harm that is objectively unjustified, and this raises the question of who should suffer this harm. In a society where it is assumed that everyone has an equally valuable choice set, with gambles that they can reasonably avoid, there is nothing unjust about asking pedestrians to absorb this harm, especially if independent moral reasons speak in favour of such an allocation.⁵³

In sum, we can say that a commitment to holding people responsible for the consequences of their choices speaks in favour of foreseeable risk-imposers' liability to defensive harm if, and only if, the risk-imposers live in a society where:

1. there is substantive equality of opportunity;⁵⁴
2. the gambles that agents take by engaging in foreseeable risk-imposing activities are reasonably avoidable;
3. it is common knowledge that foreseeable risk-imposers are expected to absorb all objectively unjustified harm that their actions might threaten;
4. there are moral reasons independent of the value of holding people responsible that speak in favour of holding foreseeable risk-imposers accountable in this way;⁵⁵
5. there is heterogeneity among agents, such that foreseeable risk impositions cannot be characterised as reciprocal.

These five conditions do not generally obtain in the world that we live in, neither with respect to driving nor with respect to other foreseeable risk-imposing activities that we permissibly engage in. It follows that for the world that we live in, the Responsibility Account generally fails to provide a justification in favour of shifting the costs

⁵³ I thank an anonymous reviewer for pressing me on this issue.

⁵⁴ As mentioned in footnote 49, it can sometimes be fairer to burden a foreseeable risk-imposer with the harm that she threatens even in the absence of equal opportunity, namely when doing so helps equalise opportunities between the foreseeable risk-imposer and her potential victim. But what makes it fairer in such cases to burden the risk-imposer is not related to her responsibility for the harmful outcome. See also footnote 55 below. I thank Alex Voorhoeve for pressing me on these issues.

⁵⁵ Such reasons (which might for example be welfarist in nature; see footnote 52) might speak in favour of holding foreseeable risk-imposers accountable even in the absence of equal opportunity and the reasonable avoidability of gambles. Having said that, if we, as a society, justifiably decide to heed these reasons and hold foreseeable risk-imposers accountable even in an imperfectly just society, what justifies our decision are the relevant *independent* reasons, and *not* the luck egalitarian thought that it is not unjust if people are worse off as a result of their voluntary choices.

of foreseeably risk-imposing activities back on the agents who chose to engage in these activities in a morally permissible manner.

V. CONCLUSION

In this paper, I have critically examined the Moral Responsibility Account of Liability to Defensive Harm. I have argued that while we can make sense of the notion of a *foreseeably risk-imposing activity* that the account relies on, the account's justificatory rationale has force only in the presence of demanding background conditions. More precisely, an agent's choice to engage in a foreseeably risk-imposing activity can successfully ground her subsequent liability to defensive harm only in a society with *substantive equality of opportunity* in which engaging in a foreseeably risk-imposing activity represents a *reasonably avoidable gamble*.

Moreover, if what I have argued is correct, then we cannot appeal to the value of holding people responsible for the consequences of their choices to settle the question whether foreseeable risk-imposers *should* be held responsible for whatever objectively unjustified harm that they might threaten. To decide what consequences should attach to which choices, we need to consult our 'substantive moral convictions that are independent of our commitment to responsibility itself.'⁵⁶

As I see it, even if our moral convictions that are independent of our commitment to responsibility lead us to conclude that foreseeable risk-imposers should be asked to internalise whatever harm that they might threaten, this would not provide us with a decisive reason in favour of accepting the Responsibility Account. As this paper has attempted to show, the Responsibility Account's appeal to responsibility has justificatory force only if demanding background conditions obtain. But I believe that many of us—and I count myself among this group—share the intuition that liability to defensive harm is a much more *robust* notion, in the sense that the factors which render an agent liable to defensive harm are not very dependent on societal background conditions. For those among us who consider such robustness a *desideratum* for a liability account, it is noteworthy that so-called culpability accounts of liability to

⁵⁶ Olsaretti, 'Responsibility and the Consequences of Choice', p. 167.

defensive harm satisfy it.⁵⁷ On culpability accounts, the key feature that grounds an agent's liability to defensive harm is the agent's culpability with respect to a threat of objectively unjustified harm. If an agent acts culpably, the agent by definition chooses an action she could reasonably have avoided—in fact, she chooses an action for which it was epistemically accessible to her that she *ought* to have avoided it. When the agent's trespass results in a threat of harm that would not have occurred had the agent acted permissibly, it seems plausible to think that the agent lacks a reasonable complaint against being burdened with the costs of her trespass. Importantly, this holds true no matter how distributively just or unjust a society the trespassing agent lives in.

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/>.

REFERENCES

- Arneson, R. (1989). Equality and Equal Opportunity for Welfare. *Philosophical Studies*, 56:77–93.
- Cohen, G. A. (2011). *On the Currency of Egalitarian Justice and Other Essays in Political Philosophy*. Princeton University Press, Princeton, NJ.
- Dworkin, R. (1981). What is Equality? Part 2: Equality of Resources. *Philosophy and Public Affairs*, 10:283–345.
- Ferzan, K. K. (2012). Culpable Aggression: The Basis for Moral Liability to Defensive Killing. *Ohio State Journal of Criminal Law*, 9:668–697.

⁵⁷ For a defence of such an account, see, e.g., Kimberly Ferzan, 'Culpable Aggression: The Basis for Moral Liability to Defensive Killing', *Ohio State Journal of Criminal Law* 9 (2012): pp. 668–697. See also Jeff McMahan, 'Innocence, Self-Defense and Killing in War', *Journal of Political Philosophy* 2(3) (1994): pp. 193–221, esp. pp. 201–204.

- Fletcher, G. (1972). Fairness and Utility in Tort Theory. *Harvard Law Review*, 85:537–573.
- Gardner, J. (2001). Obligations and Outcomes in the Law of Torts. In Cane, P. and Gardner, J., editors, *Relating to Responsibility. Essays for Tony Honoré on His Eightieth Birthday*, pages 111–144. Hart Publishing, Oxford.
- Gordon-Solmon, K. (2018). What Makes a Person Liable to Defensive Harm? *Philosophy & Phenomenological Research*, 97(3):543–567.
- Hájek, A. (2007). The Reference Class Problem Is Your Problem Too. *Synthese*, 156:563–585.
- Hart, H. L. A. and Honoré, T. (1985). *Causation in the Law*. Oxford University Press, Oxford and New York, 2 edition.
- Honoré, T. (1997). The Morality of Tort Law—Questions and Answers. In Owen, D. G., editor, *Philosophical Foundations of Tort Law*, pages 73–95. Clarendon Press, Oxford.
- Honoré, T. (1999). Responsibility and Luck. The Moral Basis of Strict Liability. In *Responsibility and Fault*, pages 14–40. Bloomsbury Publishing, London.
- Kelley, P. (2001). Restating Duty, Breach, and Proximate Cause in Negligence Law: Descriptive Theory and the Rule of Law. *Vanderbilt Law Review*, 54(3):1039–1070.
- Lazar, S. (2009). Responsibility, Risk, and Killing in Self-Defense. *Ethics*, 119(4):699–728.
- Lippert-Rasmussen, K. (2016). *Luck Egalitarianism*. Bloomsbury Publishing, London.
- McMahan, J. (1994). Innocence, Self-Defense and Killing in War. *Journal of Political Philosophy*, 2(3):193–221.
- McMahan, J. (2005). The Basis of Moral Liability to Defensive Killing. *Philosophical Issues*, 15(1):386–405.
- McMahan, J. (2009a). *Killing in War*. Oxford University Press, New York, NY.
- McMahan, J. (2009b). Self-Defense Against Morally Innocent Threats. In Robinson, P. H., Garvey, S. P., and Ferzan, K. K., editors, *Criminal Law Conversations*, pages 385–394. Oxford University Press, New York.
- McMahan, J. (2011). Who is Morally Liable to be Killed in War? *Analysis*, 71(3):544–559.
- Oberdiek, J. (2017). *Imposing Risk. A Normative Framework*. Oxford University Press, Oxford and New York.
- Olson, S. (2009). Responsibility and the Consequences of Choice. *Proceedings of the Aristotelian Society, New Series*, 109:165–188.
- Otsuka, M. (1994). Killing the Innocent in Self-Defense. *Philosophy and Public Affairs*, 23(1):74–94.
- Otsuka, M. (2016). The Moral-Responsibility Account of Liability to Defensive Killing. In Coons, C. and Weber, M., editors, *The Ethics of Self-Defense*, pages 51–68. Oxford University Press, Oxford and New York.
- Owen, D. (2009). Figuring Foreseeability. *Wake Forest Law Review*, 44:1277–1307.
- Parfit, D. (2011). *On What Matters. Volume I*. Oxford University Press, Oxford and New York.
- Quong, J. (2009). Killing in Self-Defense. *Ethics*, 119(3):507–537.
- Segall, S. (2010). *Health, Luck, and Justice*. Princeton University Press, Princeton.
- Stemplowska, Z. (2009). Making Justice Sensitive to Responsibility. *Political Studies*, 57:237–259.
- Stemplowska, Z. (2013). Rescuing Luck Egalitarianism. *Journal of Social Philosophy*, 44(4):402–419.
- Thomson, J. J. (1990). *The Realm of Rights*. Harvard University Press, Cambridge MA.
- Voorhoeve, A. (2018). May a Government Mandate More Comprehensive Health Insurance than Citizens Want for Themselves? *Oxford Studies in Political Philosophy*, 4:167–191.
- Zipursky, B. (2009). Foreseeability in Breach, Duty, and Proximate Cause. *Wake Forest Law Review*, 44:1247–1275.

Department of Philosophy, Logic and Scientific Method
 London School of Economics and Political Science, Houghton Street, London, WC2A 2AE, UK
 E-mail: s.burri@lse.ac.uk

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