



Justifying Harm

Author(s): David Rodin

Reviewed work(s):

Source: *Ethics*, Vol. 122, No. 1, Symposium on Jeff McMahan's Killing in War (October 2011), pp. 74-110

Published by: [The University of Chicago Press](#)

Stable URL: <http://www.jstor.org/stable/10.1086/662295>

Accessed: 08/12/2011 20:20

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*.

Justifying Harm*

David Rodin

In this article, I develop a general explanatory model of the liability and lesser evil justifications of harm. Despite their respective provenance in consequentialist and deontological ethics, both justifications are, at root, rich forms of the proportionality relationship between a shared set of underlying normative variables. The nature of the proportionality relationship, and the conditions under which it operates, differ between the two forms of justification. The article explores these differences in detail and the implications they have for the justification of self-defense and war.

In what circumstances and for what reasons is it permissible to harm one person in order to avert (or in the course of averting) harm to others? This question lies at the heart of some of the most contested issues in civil politics, criminal law, and war. There are two classical and competing ways of addressing the question. The first explains the permissibility of harming one to avert harm to others on the basis of a specific liability (understood as the absence or suspension of a right against harm) on the part of the person harmed. The second approach explains the permissibility of harming one to avert harm from others as a preference for the lesser evil. Consider these familiar examples:

Defense: a man happens across a villain attacking an innocent victim. The only way to save the victim's life is to kill the villain.

Nondefensive rescue: a man rushes to save an infant teetering at the edge of a precipice. In doing so, he knocks to the ground

* I have benefited enormously from several iterations of detailed comments by Seth Lazar and Jeff McMahan, and from substantial and constructive comments by two referees. Like other articles in the symposium, this article was first presented at the annual meeting of the Oxford Institute for Ethics, Law and Armed Conflict (ELAC). I am grateful for comments received there (particularly from my respondent Jonathan Quong) and at the ELAC prepublication reading group.

an innocent bystander, causing painful temporary bruising to his ribs.

Both cases are examples of justified harming, but there are significant and familiar differences between the forms of justification that underlie them. In the first case, the villain is harmed, but he is not wronged in the sense that none of his rights are infringed or violated. Neither he nor his estate is owed apology or compensation. The same is not true in the rescue case. Here the bystander is not liable to be harmed. His rights have been infringed, albeit justifiably, by the rescuer, and as a consequence he is owed some apology or compensation.

In both cases, justification depends upon a comparative assessment that significant goods can only be realized at the cost of morally acceptable harms. But what counts as a relevant good or harm for this judgment and the way in which they are assessed are different in each case. For example, if five or even fifty villains attacked the victim, they would still be liable to be killed in the defense case. Harms to multiple aggressors are not aggregated for the purposes of assessing liability to defensive harm, but harms to affected parties clearly are aggregated in lesser evil justifications like the rescue case. Similarly, if the villain were a famous surgeon, on whom the lives of five other innocent persons depend, this would not be relevant to his liability to be killed in defense. But if the bystander in the rescue case were a surgeon and the bruising would prevent him from doing his work, then this is clearly relevant to whether inflicting the harm on him is justified as the lesser evil.

Moreover, liability and lesser evil approaches to justification sometimes yield conflicting assessments. In the rescue case, it is justifiable to harm the bystander as a lesser evil, even though he has no liability to be harmed. In the modified defense case in which the villain is a surgeon, the villain is liable to be killed in defense even though killing him would not be the lesser evil.

What underlies the lesser evil and liability justifications for harm? What explains their manifest differences and the way they interact? We may begin with a very general characterization that will be developed during the course of this article. The distinctive feature of liability justifications of harm is that they concern a localized comparison between the normative status of the agency of persons in a situation of conflict. In the defense case, the villain is responsible for an unjust attack on the victim, whereas the defender is not so responsible. It is this asymmetry in their respective agency that explains why the villain is liable to be killed by the defender and not the other way around (even though both may constitute a threat to the life of the other). Lesser evil justifications, on the other hand, concern a

more generalized assessment of the comparative value of differing outcome states of affairs of the world at large. Consequences of action that would not be relevant for the purposes of assessing liability (such as the effects on multiple aggressors or those whose welfare is linked to the aggressor) are clearly relevant to assessment of the lesser evil.

In this article, I seek to develop a general explanatory account of these two differing ways of justifying harm. My conclusion will be that despite important differences between them, and despite their respective provenance in deontological and consequentialist ethics, both the liability and the lesser evil justification of harm are, at root, rich forms of the proportionality relationship between a substantially shared set of underlying normative factors. The nature of the proportionality relationship and the conditions under which it operates differ between the two forms of justification, and this gives rise to many of the practical differences between them. But as will become clear, lesser evil and liability justifications are deeply connected and interpenetrated. This raises the important question of how they interact to yield all-things-considered judgments of permissibility. I address aspects of this question throughout the article and particularly in the final section. I hope that, in addition to its theoretical interest, this account will provide a useful framework for practical decision making in the many situations where averting harm for some requires inflicting harm on others.

I proceed as follows. First, building on McMahan's account of proportionality, I argue that liability to defensive harm is in essence an issue of proportionality. Second, I identify and explain fourteen normative factors that determine whether a person is liable to defensive harm. Proportionality for liability consists in a relationship between these factors. Third, I argue that lesser evil justification emerges out of a proportionality relationship between these same fourteen factors.

Lesser evil and liability justifications of harm differ, however, in the way they treat these considerations for the purposes of proportionality. There are dramatic differences as to which harms and benefits are relevant to justification, particularly in cases in which harm-producing action will also produce benefits. In addition, certain considerations function as necessary conditions for justification, while others function as sufficient conditions for justification.

Finally, some considerations function as a simple threshold for liability, whereas others make a continuous contribution to liability. I argue that this difference is related to the nature of the good that is preserved in defensive action. Defense of noncompensable goods, like human life, allows minimum thresholds for liability. Defending compensable goods like property requires a continuum approach to the

determinants of liability. I suggest that the same distinction can explain our extreme reluctance to allow lesser evil justification for intentionally killing or inflicting grievous bodily harm. It further suggests that a more restrictive approach to collateral damage in war may be required.

I. PROPORTIONALITY

Jeff McMahan's *Killing in War* makes fundamental contributions to our understanding of the liability and the lesser evil justifications of defensive harm.¹ Throughout this article, I will respond to and build on the argument of his book to further explore these two forms of justification.

A useful starting point is McMahan's innovative account of proportionality. Proportionality is central to liability justification, but its role is puzzling. Why should liability to harm be subject to a constraint of proportionality? And what precisely does proportionality require? McMahan's treatment helps to resolve these puzzles.

First, McMahan argues that liability is intrinsically linked to the achievement of some further good or goal: "The goal is internal to the liability, in the sense that there is no liability except in relation to some good that can be achieved by harming a person."² This conception of liability to harm as instrumental (or intrinsically linked to realizing some good) fuses at the root the deontological underpinnings of liability with the seemingly consequentialist requirements of necessity and proportionality. Rights function to protect persons from being used simply as means to the ends of others. But if the possession of rights can be conditional on the observance of relevant moral requirements, and in particular on respecting the rights of others, then it is natural to think that rights can be forfeited in precisely this way: by transgressing a relevant moral requirement, a person can become liable to be harmed as a means to preventing or remedying that very transgression. Necessity and proportionality therefore follow as necessary components of liability.

Second, McMahan argues that the demands of proportionality are sensitive to two considerations: intentionality on the part of the defender, and liability on the part of the person who suffers the defensive harm. Liability is in turn analyzed as a product of two further

1. Liability to defensive harm is of course only one species of a broader category of liability justifications that includes liability to punitive harm and liability to harm for the purposes of redress. For the remainder of this article, I focus on the liability justification for defensive harm, though as I will bring out in the last section of this article, there are important connections between liability to defensive harm and liability to redress harm.

2. Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), 8.

considerations: an objectively unjustified threat of harm, and moral responsibility for that harm on the part of the potentially liable person.³ This leads McMahan to posit the existence of four distinct forms of proportionality judgment:

1. Acts that intentionally harm those who are potentially liable to be harmed (proportionality as traditionally understood in personal self-defense).
2. Acts that unintentionally but foreseeably harm those who are potentially liable.
3. Acts that intentionally harm those who are not liable.
4. Acts that unintentionally but foreseeably harm those who are not liable to be harmed (proportionality of side-effect harms as classically conceived in *jus in bello*).⁴

He refers to 1 and 2 as ‘narrow proportionality’. This is the form of proportionality relevant to what I call liability justifications for defensive harming. He refers to 3 and 4 as ‘wide proportionality’, which is the form of proportionality that is relevant to lesser evil justifications for harm.

This fourfold taxonomy is a genuine innovation. Yet even this enriched understanding does not appropriately capture the role that proportionality plays in the justification of defensive harm. McMahan’s classification is at once overly simple and overly complex. It is overly simple because it makes certain determinants of proportionality constitutive of the taxonomy, while ignoring others. For example, McMahan’s taxonomy rightly highlights the fact that proportionality is sensitive to the intention of the person who inflicts defensive harm: directly intended harm is harder to justify and therefore subject to a more demanding proportionality constraint than harm that is foreseen but unintended. But as McMahan clearly recognizes, if the intention of the defender is relevant to the proportionality judgment, then so too is the intention of the aggressor. This should be explicitly recognized in our account of proportionality.

Similarly, if the distinction between intention and foresight is relevant to proportionality, then it seems likely that the distinction between doing and allowing will also be relevant. Harm brought about through something we do (‘positive agency’ in Quinn’s terminology)⁵

3. For McMahan, these two considerations are both necessary and sufficient for liability to harm: “The criterion of liability to attack . . . is moral responsibility for an objectively unjustified threat of harm” (ibid., 35). As will become clear, I am doubtful that these two conditions are sufficient for liability to defensive harm.

4. Ibid., 20.

5. Warren Quinn, “Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing,” *Philosophical Review* 98 (1989): 287–312.

is, other things being equal, more difficult to justify than harm brought about by something we allow to happen (negative agency). As with the distinction between intention and foresight, the distinction between doing and allowing will have relevance on both sides of the proportionality relationship; that is to say, it will be relevant to the harm inflicted in the course of defensive action and to the harm the defensive action seeks to avert.

One might conclude that we need to expand the species of proportionality from four to eight, in order to reflect the interaction of the doctrine of doing and allowing with the considerations of intention. But in fact, as will quickly emerge, the considerations of intention versus foresight and of doing versus allowing are but two among a considerable number of normative factors that affect proportionality. Rather than create proliferating species corresponding to each determinant of proportionality, we should seek a unified account that integrates all relevant factors and provides some coherent way of relating them. That is what I will attempt in this article. This is the sense in which McMahan's analysis is overly complex. Rather than four basic species of proportionality, there are just two basic forms corresponding to McMahan's distinction between narrow and wide proportionality (though the internal structure of each is more complex than McMahan's treatment suggests).

We can be even more parsimonious in our analysis. Traditional theories of defensive rights view liability and proportionality as distinct moral operators in the justification of harm. We ask first whether some person is liable and then inquire whether the harm is proportionate. This is reflected in the traditional view that proportionality consists in a simple comparison between the harmful and beneficial consequences of action.

But both assumptions are wrong. There is no independent condition of liability to harm separate from considerations of proportionality, and no harm can be described as proportionate without reference to the liability of the person affected. This is because a person can only be liable to a particular harm that is proportionate in the circumstances (if the harm were not proportionate, he would not be liable to it). As McMahan says, "the restrictions on liability are . . . 'internal' to liability itself."⁶ Thus, it makes no sense, say, that someone is liable to harm *simpliciter*, without specifying the harm to which he is liable. For a person to be liable to a harm, just is for that harm to be narrowly proportionate in the circumstances. Proportionality and liability, far from being independent factors, are two manifestations of the same underlying normative relations.

6. McMahan, *Killing in War*, 10

A consequence is that proportionality for liability should not be conceived as a relationship between instances of good and harm at all. Rather, it is a relationship between the normative status of the acts of agents (of which the good and harm that they produce are but one contributing factor). Understanding this will help to resolve a long-standing dispute. Does proportionality require good consequences of action to exceed bad consequences, or merely that they be roughly commensurate? It is clear that if factors such as intention and doing/allowing are relevant to proportionality, then the dispute about what proportionality requires cannot be resolved so long as one is posing the question in terms of a simple comparison between harms and goods. For whether the beneficial consequences of action are required by proportionality to exceed harmful consequences or merely be roughly commensurate with them will depend critically on the normative quality of the agency by which they are brought about.

II. FACTORS THAT DETERMINE LIABILITY TO DEFENSIVE HARM

I will identify fully fourteen factors that determine liability to defensive harm and also thereby narrow proportionality. I suggested above that liability arises out of a contrast between the agency of two persons in a situation of conflict. Liability-determining factors can therefore be divided into two groups: (A) factors that concern a normative assessment of the agency of the harm-threatening actor, and (B) factors that concern a normative assessment of the agency of the actor who inflicts harm in the course of averting the threatened harm. These factors may be summarized prior to a fuller discussion:

Factors Relevant to the Threatening Agent

1. Magnitude of the threatened harm.
2. Probability of the threatened harm occurring.
3. Responsibility for the threatened harm.
4. Justification of the threatened harm.
5. Whether the harm is brought about through doing or allowing.
6. Intention with which the threatened harm is brought about.
7. Aggravating circumstances such as preexisting duties of care.
8. Causal and temporal proximity to the threatened harm.

Factors Relevant to the Defending Agent

9. Magnitude of the defensive harm.
10. Probability of the defensive harm averting the threatened harm.
11. Intention with which the defensive harm is brought about.
12. Responsibility for the threatened harm.
13. Whether the defensive harm is brought about through doing or allowing.
14. Preexisting duties of care.

In order to facilitate exposition of these factors, let us consider a generic form of cases in which liability to defensive harm may be at issue. We have a relationship between two agents, A and D, where A's action will inflict threatened harm on D unless D acts in a way that will inflict defensive harm on A. In standard cases of self-defense, of course, A is an aggressor and D is a defender. However, we must specify the case more generally to allow for circumstances in which A may impose harm or the risk of harm without being an aggressor or without even directly posing a threat of harm at all, for example, cases of negligently or recklessly imposing risks on others.

For the purposes of this specification, any potential harm may be considered as a threatened harm. The threatened harm need not be the first in a temporal sequence of harm, and even a harm inflicted in the course of self- or other-defense may be considered a 'threatened harm' for the purposes of moral assessment.

Defensive harm, on the other hand, is defined by its relationship to a given threatened harm. Defensive harm is the causal product either of action that does avert the threatened harm or of action that is intended by D to avert the threatened harm. Thus, the infliction of defensive harm on A may be the means by which D seeks to avert threatened harm (as in standard cases of self-defense), or it may be the foreseen or unforeseen side effect of action by which D seeks to avert threatened harm (as in standard cases of potentially justified side-effect harm).

This formulation leaves open whether threatened harm will be inflicted on D, the agent of defensive harm, or some third party. Effects on third parties introduce considerable complexities that I address in the final sections.

III. LIABILITY-DETERMINING FACTORS RELEVANT TO EVALUATION OF THE AGENCY OF A

Whether A is liable to defensive harm by D will depend on at least the following factors:

1. *The Magnitude of the Threatened Harm*

A's liability to defensive harm varies with the magnitude of the threatened harm for which he is responsible. It is justifiable to inflict greater defensive harm to foil a murderer than to foil a pickpocket. This is because, other things being equal, it is worse to engage in action that brings about greater compared with lesser harm. Because liability to harm arises out of a localized asymmetry in the normative status of agents, not all harmful or beneficial consequences of A's action may contribute to the magnitude of the threatened harm. This raises difficult questions, to which we will return below, about how to demarcate the boundary of the threatened harm when action that causes the threatened harm will also generate benefits for A, D, or third parties.

2. *The Probability of Threatened Harm Occurring If Defensive Harm Is Not Inflicted*

In many cases, the probability of the threatened harm occurring, in absence of the action that generates defensive harm, will also have an effect on A's liability to harm. For example, if I would be justified in inflicting \$100 of damage to your faulty automatic watering system which was otherwise certain to inflict \$100 of water damage to my new carpet, it seems plausible that I would be justified in inflicting a lesser degree of damage if the probability of damage was only 10 percent. However, in other cases, particularly those concerning threats to life and grievous bodily harm, probability of harm does not seem to play an analogous role. For example, it would seem permissible to use lethal force to prevent A from playing Russian roulette with D even if only one of the gun's ten chambers was loaded. Why this should be so is an important question to which we will return below.

3. *The Responsibility of A for the Threatened Harm*

One of the central themes of McMahan's book is that for A to be liable to defensive harm, it must be the case that A has moral responsibility for the threatened harm to which defensive measures are a remedy. This requirement is clearly central to the protection afforded by rights: rights against harm can be alienated only on the basis of some aspect of the right-holder's responsible agency.

There are a number of ways in which a harm may fail to be appropriately attributed to A's agency. Each will affect liability:

3.1. A is an *innocent bystander* in the sense that he has no agency involvement in bringing about the threatened harm at all. Almost all authors agree that an innocent bystander has no liability to be

harmed in defense, even if this were the only way to avert significant threatened harm.

3.2. A has an *excuse* for action which results in the threatened harm. McMahan largely follows George Fletcher's classic account of justification and excuse, maintaining that while justifications exempt an agent from both liability to punishment and liability to defensive harm, excuses exempt an agent only from liability to punishment and leave the defensive rights of others undisturbed.⁷ Both aspects of this claim are open to dispute, however. As I shall argue, some excuses may exempt agents from liability to defensive harm, whereas some full justifications may not. In particular, excuses can be of two different kinds:

3.2.1. *Agency-diminishing excuses.*—It is indeed true that many common excuses do not exclude liability to defensive harm, even when they provide full exculpation for the purposes of punishment. These include:

1. Excuses which deny that the harm producing action was intentional, or that it was intentional under the proscribed definition. Examples include action that is inadvertent or made under a reasonably mistaken interpretation of the facts.
2. Excuses which concede that the harm-producing action was intentional but deny that it was voluntary: for example, duress, necessity, and provocation.
3. Excuses which claim that the person lacked the capacity for full deliberative agency. Infancy, insanity, and involuntary intoxication are excuses of this form.⁸

Though these excusing conditions diminish the attribution of the act to the agent, they do not defeat it entirely. This minimal attribution of responsibility is often sufficient to ground potential liability to defensive harm.⁹ For example, psychotic aggressors, aggressors who are acting under conditions of extreme duress, and child aggressors are normally considered to be liable to be killed in self-defense, even if one concedes that the excuses are fully exculpating for the purposes of punishment. McMahan argues further that even though excused actors can be liable to defensive harm as long as they have minimal

7. George Fletcher, *Rethinking Criminal Law* (Boston: Little, Brown, 1977), 760.

8. McMahan employs a more fine-grained analysis of overlapping categories: Partially Excused Threats, Excused Threats, and Innocent Threats, which may all in some respects be morally responsible. They are contrasted with Non-responsible Threats (*Killing in War*, 161ff.).

9. Seth Lazar helpfully refers to persons who possess fully exculpating excuses of this form as possessing "agent responsibility" ("Responsibility, Risk, and Killing in Self-Defense," *Ethics* 119 [2009]: 699–728).

responsibility for an unjustified threatened harm, nonetheless, the fact of their excuse can affect the proportionality of defensive action. For example, he argues that one may be required to use lesser force or accept a higher degree of risk in using defensive force against a partially excused aggressor than one would against a fully responsible aggressor.¹⁰

3.2.2. *Agency-defeating excuses* (excusing conditions that generate what McMahan calls nonresponsible actors).—In contrast to excuses which function by demonstrating degraded agency, a further class of excuses leaves no room whatever for agency in the production of the threatened harm. The locus classicus of an agency-defeating excuse is physical compulsion, as in Robert Nozick's famous case of the man who has been thrown down a well and who will crush D unless he vaporizes him with his ray gun. Many authors, including Jeff McMahan, Michael Otsuka, and myself, have argued that persons who threaten harm with an agency-defeating excuse such as physical compulsion are indistinguishable from innocent bystanders and hence not liable to defensive force, though this remains a contested view.¹¹

4. *A's Justification for Bringing about the Threatened Harm*

As we have seen, excuses in their agency-diminishing form are consistent with liability to defensive harm, though the presence of an excuse may diminish liability to defensive harm. Justifications, on the other hand, typically do exclude liability to defensive harm completely. However, I will suggest that one form of justification (that which consists in the justified infringement of a right) is consistent with liability to defensive harm. Justifications for inflicting threatened harm may fall into four categories:

4.1. *Harms that are not proscribed by any moral or legal norm.*—Many common side-effect and externality harms are of this form. For example, a driver contributes to the harm of congestion, and a business proprietor may harm competitors through competitive pricing. These harms, though real, do not generate defensive rights because they are harms of a species that are not proscribed, and against which no one has a right.

McMahan further distinguishes justified acts from permitted acts.¹² Justified acts are said to be those for which there is a positive moral reason for their performance, whereas permitted acts are sim-

10. McMahan, *Killing in War*, 158ff.

11. Jeff McMahan, "Innocence, Self-Defense and Killing in War," *Journal of Political Philosophy* 2 (1994): 193–221. See also McMahan, *Killing in War*, chap. 4; Michael Otsuka, "Killing the Innocent in Self-Defense," *Philosophy and Public Affairs* 23 (1994): 74–94; David Rodin, *War and Self-Defense* (Oxford: Oxford University Press, 2002), 79ff.

12. McMahan, *Killing in War*, 43.

ply those that are not proscribed. I do not find the distinction helpful because most permitted acts that are actually performed will have some positive moral reason for their performance, even if it is a very weak one. For example, the fact that a permissible act contributes to economic activity, brings me pleasure, or satisfies a desire provides *some* moral reason for its performance. Most of the acts that McMahan describes as permitted are in fact weakly justified under his use of this term.

The category of nonproscribed harms may hold the solution to the trouble McMahan has with his case of the conscientious driver. In this case, a “freak event” causes a driver who has taken all reasonable care in the maintenance and control of his vehicle to veer suddenly toward a pedestrian who can save his life only by blowing up the conscientious driver’s car and the driver with it.¹³ McMahan believes that the driver’s responsibility for imposing a small risk of unjustified harm on the pedestrian is sufficient to ground liability to lethal defensive harm.

However, this is an intuitively uncomfortable result. It can be avoided if we distinguish between responsibility for imposing the risk of harm, and responsibility for the harm itself. The driver is responsible for imposing a small risk of injury on the pedestrian, but imposing that risk was not proscribed given that he had fulfilled all his obligations to minimize the risk. The pedestrian has no right not to be exposed to such a risk. Striking the pedestrian with a car clearly is proscribed, but that is arguably not an action for which the driver is responsible. Much depends on the nature of the freak event that causes the accident. If the accident resulted from mechanical failure such as failed brakes or seized steering, then it would be hard to describe the strike as an action of the driver at all. The driver has now become a passenger in an out-of-control vehicle, and his situation seems comparable to cases of physical compulsion like the falling fat man.

But our analysis may be different if the freak event was instead excusable inadvertence on the part of the driver, for example, if he momentarily took his eyes off the road. In such a case, the driver has at best an agency-diminishing excuse for striking the pedestrian that is indeed consistent with liability, and intuitively it seems more plausible that he may be liable to defensive harm.

4.2. *Harms against which D has alienated his right.*—A may be justified in inflicting threatened harm on D because D has, through his own responsible agency, alienated his right not to have the threatened harm inflicted on him. The right may have been voluntarily alienated

13. *Ibid.*, 165ff.

as a result of consent, sale, disposal, waiver, or the like. Alternatively, the right may have been involuntarily alienated through liability assumed as a consequence of responsibility for some transgression. Three forms of liability to harm may result from transgression: liability to defensive harm, liability to punitive harm, and liability to redress harm. As in the case of harms that are not proscribed, threatened harm to D against which D has alienated his right does not generate liability to defensive harm.

4.3. *Harm that is all-things-considered justified as the lesser evil.*—A may be justified, all things considered, in inflicting threatened harm on D even though D has a right that A not inflict the harm, because it is necessary to avert some substantially greater moral evil. In such a case, D's right is not forfeited, but neither is it violated; rather, we say that it is justifiably infringed. An example that McMahan discusses is a bomber pilot who is about to inflict foreseen but unintended necessary harm on innocent civilians in the service of a just cause.¹⁴ Because the civilians are not liable to be killed, the collateral harm inflicted on them would infringe their rights. McMahan believes that all-things-considered justification defeats liability to defensive force and that the pilot is therefore not liable to defensive harm inflicted by the civilians.¹⁵ After all, it seems odd that you could lose significant rights against harm simply for doing what morality all-things-considered permits you (or even requires you) to do.¹⁶

There are, however, reasons to be skeptical of this view. Note that although lesser evil justification defeats liability to punishment, it does not defeat many other significant forms of liability, for example, the liability to pay back debts or the liability to compensate for harm. Suppose that I am morally required to write a check to Oxfam, and that doing so puts me into overdraft. The bank will not be very impressed if I say: "Obviously I am not liable to repay this debt because my donation was all-things-considered justified, and justification defeats liability." Similarly, in the nondefensive rescue case considered in the introduction, affirming that the rescuer was all-things-considered justified in bruising the ribs of the bystander is consistent with believing that he is liable to make good the harm he has inflicted in some way—perhaps by apologizing, tending to him in the hospital, or making a financial contribution.

14. In the final section of this article, I will argue that the conditions for justifying the collateral harming of civilians are considerably more restrictive than is typically accepted. The permissibility of collaterally killing the civilians is here accepted for the sake of argument.

15. *Ibid.*, 41ff.

16. Jeff McMahan, "The Basis of Moral Liability to Defensive Killing," *Philosophical Issues* 15 (2005): 386–405, 399.

I believe that liability to defensive harm may be similarly compatible with an all-things-considered lesser evil justification for harming. The bomber pilots are objectively justified in inflicting incidental harm on the civilians. Their actions are not wrong, all things considered, but still they *wrong* the civilians, in the sense that they infringe their rights. Just as the bank manager can reasonably ask: “What business is it of mine that you are morally required to give to Oxfam? You have borrowed my money and you are liable to pay it back,” so the civilians can argue “What business is it of mine that you are morally required to undertake this bombing mission? You are infringing our rights and you can be liable to defensive measures required to uphold those rights.” In both cases, the thought seems linked to the principle that one is required to bear the costs of one’s own action even when one responds appropriately to objective moral reasons.¹⁷

As was suggested in the introduction, liability justifications emerge out of a narrow and localized comparison between the normative status of the agency of persons in a situation of conflict. Many of the broader value considerations that are relevant to lesser evil justification are irrelevant to liability justification. An underlying reason for this may be that many of the rights at issue in liability justifications are reciprocal in character. Arguably, I have the right that you not kill me in part because (and to the extent that I do) respect your reciprocal right I not kill you.¹⁸ If this interpersonal reciprocity is what underlies rights like the right to life, then it is easy to understand why A’s broader justification for infringing D’s rights may be irrelevant to his liability to be defensively harmed: justified infringements of rights also breach reciprocity.¹⁹

The view that justified infringement of rights can potentially ground liability to defensive harm is more plausible if one remembers that liability to harm consists simply in the absence of a right against being so harmed. It does not in itself determine the broader permissibility of inflicting the harm. It is indeed possible that the pilots in McMahan’s bomber case are liable to be killed in self-defense by the civilians, but that the civilians all-things-considered ought not to kill them, because of the broader goods at stake in the mission. Just as lesser evil considerations may sometimes justify inflicting harm on someone who is not liable, they may sometimes prohibit inflicting harm on someone who is liable.

17. This is sensitively discussed by McMahan in *Killing in War*, 47ff.

18. See Rodin, *War and Self-Defense*, chap. 4.

19. See Sec. V.C.i below for an additional argument for this conclusion.

5. *Whether A Brings about the Threatened Harm through Doing or Allowing*

Like the distinction between intention and foresight which McMahan makes central to his analysis of proportionality, the distinction between doing and allowing plays a role in determining liability to defensive harm and hence also to proportionality. For example, it is plausible that, other things being equal, D would be entitled to inflict greater defensive harm on A if A had pushed a trolley down a track on which D was standing, than if A had simply failed to stop a runaway trolley headed for D. The *ceteris paribus* clause is important because, as we will see, threatened harms intentionally brought about through negative agency can ground defensive rights as strong as those bought about through positive agency.

6. *A's Intention in Bringing about the Threatened Harm*

Intention plays a role in determining liability, comparable in importance to the role of responsibility. It has long been recognized that the intention of the defending agent, D, plays a critical role in the justification of defensive action. But the intention of the harm-threatening agent, A, is equally important.

6.1. *Directly intended unjustified harms.*—Such harms can clearly ground liability to defensive harm. Indeed, the direct intention to produce harm plays a particularly decisive role, since it can ground liability to defensive harm even if the threatened harm is subject to one of the other mitigating conditions considered here. For example, directly intended threatened harm can generate liability to defensive harm even if it was (a) brought about as a consequence of allowing something to happen rather than causing it through some positive agency, (b) partially brought about through the role of wrongful intervening action of others, or (c) brought about through temporally or causally distant or peculiar mechanisms. Moreover, this seems to remain true even if there is highly diminished responsibility for the directly intended harmful action, for example, if the harm-producing action was excusable due to duress.

6.2. *Unintended harmful consequences of permissible action.*—I argued above that it is possible that justified infringements of rights, for example, justified collateral harm in war, can generate liability to defensive harm. However, as the doctrine of double effect reminds us, harms that are brought about as the foreseen but unintended side effect of otherwise permitted action are easier to justify than harms that are directly intended. This may be as true at the level of the interpersonal relations that underlie rights and liability as it is at the broader level of lesser evil considerations. It is arguably a greater transgression to have one's rights intentionally violated or infringed

than it is to have those same rights infringed or violated as the foreseen side effect of otherwise permissible action. If that is right, then we should expect that foreseen but unintended threatened harms will generate lesser liability to defensive harm than directly intended threatened harm. This will not invariably be the case, however, as we must also be attentive to whether unintended harms were negligently or recklessly brought about.

6.3. *Unintended harmful consequences of impermissible action (both foreseen and unforeseen).*—These generally do ground liability to defensive harm. There may, however, be limits to the attribution of liability on the basis of unintended consequences of impermissible action when the causal relationship between the action and the harm is distant or peculiar (see point 8 below).

7. *Aggravating Conditions for Culpability*

We saw in point 3 that if A has an agency-diminishing excuse, he can be liable to defensive harm even if he is not culpable for inflicting the threatened harm. However, there is no doubt that if A is not merely minimally responsible for the threatened harm but also culpable, this can increase his liability to defensive harm. If that is right, then the presence of aggravating conditions for culpability will be relevant to liability.

A particularly important form of aggravating condition arises from preexisting duties of care. If A has a duty of care toward the person who will bear the threatened harm, then his harm-producing action is particularly egregious. This may partially explain some people's intuition that "battered wives" are permitted to engage in defensive acts that would not be permitted outside the context of a marriage.

8. *Temporal and Causal Proximity*

Traditional accounts of defensive rights limit liability to defensive harm to circumstances in which there is temporal and causal immediacy between the action of A and the threatened harm.

8.1. Imminence is the traditional requirement of temporal proximity for liability to defensive harm. This requirement rules out "preventive" action that inflicts defensive harm significantly prior to the infliction of threatened harm. It also rules out inflicting defensive harm on a "past aggressor." For example, suppose D suffers from a life-threatening injury culpably inflicted by A one year ago. If the only way for D to save his life was to kill A and harvest his organs for transplant, it is not clear that A would be liable to be killed.

8.2. Liability to defensive harm also seems restricted to cases in which there is sufficient causal immediacy between A's action and the

infliction of threatened harm. For example, if A were a cutler who made a knife used by a third party to threaten the life of D, most people would not believe that A is liable to be killed even if this were the only way to avert the attack.²⁰

There seems little doubt that temporal and causal proximity can affect liability, but the crucial question is whether they are relevant solely because of their contribution to A's responsibility for the threatened harm, or whether they play an independent role. This question is of great importance to whether noncombatants can ever be permissible targets in war. McMahan argues that in certain rare circumstances, noncombatants can be liable to be directly attacked if they are morally responsible for the harms that constitute the just cause.²¹ However, if causal proximity plays a role in determining liability, then this may help to resist this troubling conclusion, since most noncombatants do not play a causally proximate role in the unjust harms of war even if they have some degree of moral responsibility.

McMahan introduces a case to support his contention that moral responsibility for unjustified harm is sufficient for liability, even where there is a tenuous causal link between A and the threatened harm. In this case, a corrupt sheriff dupes and coerces a simple farmhand into killing the local mayor.²² McMahan argues that the sheriff is liable to be killed in preference to the farmhand, even though the farmhand but not the sheriff is posing the proximate threat to the mayor. In a similar way, suggests McMahan, noncombatants can be liable to be killed in preference to combatants if they have greater responsibility for an unjust war.

However, different cases suggest otherwise. Suppose a criminally insane psychopath who has been improperly released from the hospital because of a financial crisis in the health system threatens your life. The financial crisis is a direct result of the minister of health's criminally fraudulent mismanagement of finances in full knowledge that his actions would endanger the public. Suppose you could save your life either by killing the psychopath or by killing the minister of health (either by using him as a human shield or by riding roughshod over him in escape). It would be permissible to kill the psychopath in self-defense, despite his highly diminished responsibility, but it would seem impermissible to kill the minister even though he has greater moral responsibility for the existence of the unjust threat.

20. As noted above, the mitigating effect of causal distance can be defeated by direct intention to harm, for example, if the cutler had fashioned a special blade with the direct intention to penetrate D's protective clothing.

21. McMahan, *Killing in War*, chap. 5, esp. 221ff.

22. *Ibid.*, 205–8.

Cases like this give us reason to be cautious about McMahan's argument, but ultimately the issue is unlikely to be settled at the level of intuitive assessment of cases. What is required is a good theoretical explanation for why causal and temporal proximity affect liability independent of responsibility. This is not easy to provide. One potential explanation invokes Warren Quinn's distinction between eliminative and manipulative agency. Defensive harm inflicted on persons who are causally or temporally remote from the threatened harm tend to be manipulative in nature, which may make them harder to justify (see Sec. 11.1 below).

A second potential explanation invokes a principle of intervening agency: where two or more persons share responsibility for unjust threatened harm, defensive force should be presumptively directed at the agent whose intervening action is most proximate to the threat.²³ However, in order to be plausible, this principle will have to allow for cases in which the proximate agent is an institutional or collective agent. For example, in war, defensive force can be employed against anyone within the chain of command, not merely those who fire the guns. Similarly, any member of a criminal conspiracy seems potentially liable to defensive force. What seems to make the difference in these cases is a particularly strong form of shared intention to bring about harm, often existing in a formal or institutional context.

Interestingly, this idea of unified intent may explain the difference between the sheriff and the minister cases. The sheriff has the capacity to deputize the farmhand, and even if he is not legitimately deputized, it seems reasonable to view him as standing in a chain of command with the sheriff. The minister, on the other hand, is morally responsible for the harm inflicted by the psychopath, but there is no formal or institutional setting that unifies them in an intent to harm.

The situation of most civilians is different again. While civilians are linked to unjust soldiers through democratic institutions, they are certainly not connected to them in a chain of command. Suppose we alter McMahan's example so that the only way the mayor could defend himself was by killing a citizen of the town who had voted for the sheriff. Even if the sheriff had run on a "kill the mayor" platform and the citizen had voted for him with the direct intention of bringing about the mayor's death, it seems unlikely that the citizen would be liable to be killed. This suggests at least that participating in a democratic process that results in an unjust threat of harm is unlikely to be sufficient to yield liability to lethal force.

23. This idea is explored in David Rodin, "The Moral Inequality of Soldiers: Why *Jus in Bello* Asymmetry Is Half Right," in *Just and Unjust Warriors: The Moral and Legal Status of Soldiers*, ed. David Rodin and Henry Shue (Oxford: Oxford University Press, 2008), 50ff.

IV. LIABILITY-DETERMINING FACTORS RELEVANT TO EVALUATION OF THE AGENCY OF D

We have so far reviewed eight liability-determining factors that all concern an evaluation of the agency of A, the agent whose action causes the Defensive Harm. To complete our account of the liability justification for defensive harm we must now consider those liability-determining factors that concern an evaluation of the agency of D, the agent whose action causes that Defensive Harm:

9. *The Magnitude of the Defensive Harm (Harm That Is the Causal Product of Action That Either Does Avert the Threatened Harm, or Is Intended by D to Avert the Threatened Harm)*

Greater magnitudes of defensive harm are more difficult to justify and require a commensurate increase among factors 1–8 in order to make it proportionate. As in the case of the threatened harm, there are difficult issues, to which we will return below, about how to draw the boundary of the relevant defensive harm when action that causes the defensive harm will also generate further beneficial consequences for either A, D, or third parties.

10. *The Probability That the Harm-Producing Defensive Action Will Avert the Threatened Harm*

If D's action inflicts defensive harm on A, but there is an extremely low probability that this action will succeed in averting the threatened harm, then A's liability to that defensive harm can be diminished. At the limit, if harm-inflicting defensive action has no prospect of averting the threatened harm, then there is no liability. This is the *jus ad bellum* principle of prospect of success.

Danny Statman thinks this principle is deeply puzzling.²⁴ Surely, he argues, a woman with two bullets in her gun who faces rape by five men would be permitted to kill two of her assailants, though there is no prospect that she can avoid being raped. However, this and other group-based examples discussed by Statman do not serve his purposes well. We can see this by recalling that liability is rooted in a relationship between the agency of interacting persons. The two men whom the victim can shoot are individually liable to be killed because this measure would succeed in averting *their* rape. That this action would not prevent separate rapes by the remaining three men is irrelevant to the liability of the two.

Statman's concerns are better revealed in cases of hopeless de-

24. Daniel Statman, "On the Success Condition for Legitimate-Defense," *Ethics* 118 (2008): 659–86.

fense against a single aggressor. Suppose a victim could not hope to prevent a strong assailant from raping her, but by struggling she could break his arm. It seems peculiar to say that she would violate his rights if she did so.

But it may be that delaying the infliction of threatened harm can provide a basis for liability even if ultimate prevention is not possible (sometimes delaying harm is all we can hope for—as John Maynard Keynes famously observed, “in the long run we are all dead”). Moreover, when we say there is no prospect of success, we often mean there is a very low prospect. Harmful defensive action with a low prospect of success seems to be more justifiable in defense of life and threats to the integrity of the person such as rape than for threats to property, an idea we will return to below. If inflicting harm on A would not prevent, delay, or ameliorate the threatened harm in any way, then it is hard to see how A could be liable to the harm as a matter of defense. I suspect that those who believe that it would be permissible for the victim to break the rapist’s arm in an act of truly futile resistance would also be committed to allowing the victim to break the rapist’s arm in retribution after the fact, or to permit the victim to inject the threatening rapist with a substance that would kill him only after the crime was consummated. Neither of these later acts can be explained on the basis of liability to defensive harm.²⁵

11. *D’s Intention in Bringing about the Defensive Harm*

As McMahan correctly notes, A’s liability to defensive harm can be significantly affected by the intention with which D brings about that harm.²⁶ There are a number of possibilities:

11.1. *Inflicting defensive harm as a means to avert the threatened harm.*—In standard defense cases, D directly intends to inflict defensive harm on A as a means to averting the threatened harm. This is readily comprehensible on McMahan’s instrumental understanding of liability. If rights ordinarily protect one from being harmfully used as a means to some other person’s goals, liability (the localized suspension of those rights) consists in a moral vulnerability to be used in precisely this way.

Eliminative and opportunistic agency.—There may be limits to the extent to which one can be liable to be used as a means by others. Within the class of directly intended harms, Warren Quinn further

25. Statman’s own solution to the puzzle is to claim that futile acts of resistance do succeed in defending the victim’s honor. But the argument depends on accepting an overly narrow conception of honor that is implausibly restricted to violent resistance to offense. It implicitly denies the honor of a principled commitment to nonviolence found in the Ghandhian tradition and in aspects of the Christian tradition.

26. See, in particular, McMahan, *Killing in War*, 28–29.

distinguishes between harms that are brought about through opportunistic agency, and those that are brought about through eliminative agency. In opportunistic agency, the presence of the person harmed conduces to, or is necessary for, the success of the agent's plan or strategy (an example would be seizing a bystander and using him as a human shield against an aggressor). In eliminative agency, the victim of harm features only as a problem or obstacle to be removed (an example would be pushing a person from a bridge in order to clear a path to escape from an aggressor). Quinn suggests that harm inflicted through opportunistic agency is more difficult to justify because it involves using a person in a more objectionable way than harm inflicted through eliminative agency.²⁷

We have already seen that traditional accounts of self-defense restrict liability to cases in which there is causal and temporal immediacy between A and the threatened harm. The distinction between eliminative and opportunistic agency may in part underlie these restrictions by providing an explanation for the intuitive idea that justified defensive harm must be a case of warding off a person who is currently posing a threat.

11.2. *Inflicting defensive harm as an end in itself.*—D may directly intend to inflict defensive harm on A, not as means to averting the threatened harm, but as an end in itself. There are two further possibilities:

11.2.1.—D acts with the direct intention of inflicting the defensive harm, and the fact that his action has the side effect of averting the threatened harm is unforeseen. For example, consider a gangster who enters a bar in order to murder a rival. Unbeknownst to the gangster, his rival is about to shoot an innocent man in the bar with a concealed weapon. In this case, it would seem that the rival is not liable to be shot by the gangster, even though he would be liable to be shot by a third party acting with an intention to defend the life of the innocent man. This demonstrates that liability to defensive harm is a function of more than simply A's responsibility for an unjustified harm. It is a function also of D's intention with respect to the defensive harm. As we will see shortly, A's liability to defensive harm can also be a function of D's responsibility for the threatened harm.

11.2.2.—A more difficult case is one in which D acts with the direct intention of inflicting the defensive harm and foresees but does not intend the side effect of averting the threatened harm. This is a

27. Warren S. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Double Effect," *Philosophy and Public Affairs* 18 (1989): 334–51, at 344. This suggestion has been interestingly developed by Jonathan Quong, "Killing in Self-Defense," *Ethics* 119 (2009): 507–37.

complex case, but arguably just as we discount the foreseen but unintended collateral harms of our action for the purpose of determining liability to defensive harm, so we should discount the foreseen but unintended benefits of our action. On this reasoning, the foreseen but unintended side effect of averting the threatened harm has some weight in generating a liability to defensive harm in A, but not as much as when D directly intends the defensive harm as a means to averting the threatened harm.

11.2.3.—A final case is one in which D foresees that defensive harm will be inflicted on A as a result of action necessary to avert the threatened harm, but does not intend to inflict the defensive harm on A, either as end in itself or as a means to any further goal. This is McMahan's second species of narrow proportionality judgment. McMahan is drawing our attention to the fact that just as a person may be liable to intentional harm in the course of defense, he or she may also be liable to unintended harms. As in the case of harms that infringe rights but are justified by reason of the lesser evil, the burden of justification appears to be lower for harms that are foreseen but not intended than they are for harms that are directly intended. This will have significant implications for the ethics of war because it raises the possibility that although most noncombatants will be immune from direct targeting, they may nonetheless be liable to have certain forms of collateral harm inflicted on them.

12. *D's Responsibility for the Threatened Harm*

Liability to defensive harm is clearly sensitive to facts about A's responsibility for bringing about the threatened harm. But what is less commonly noticed is that it is also sensitive to facts about D's responsibility in bringing about the threatened harm. For example, suppose D culpably provokes A (who has diminished responsibility for his action) into attacking him, so that he may kill A and gain exoneration by means of self-defense. It seems clear that A is not liable to be killed by D, even though A bears sufficient moral responsibility for the threatened harm that D would be at liberty to kill him if a third party had provoked A into attacking D. The reason would seem to be that a right of self-defense is grounded in the normative contrast between the agency of A and D in bringing about the unjustified threat of harm, not simply in the agency of A. In cases where D has preponderant responsibility for the threatened harm or where responsibility is shared (such as in the culpable provocation case) then there is no clear normative contrast favoring D, and A's liability to defensive harm is diminished or absent.

13. *Whether D Brings about the Defensive Harm as a Result of Doing or Allowing*

Defensive harm brought about by something D does (harm brought about through positive agency) is more difficult to justify, others things being equal, than defensive harm that D allows to happen (harm brought about through negative agency). Imagine, for example, that D is standing at the edge of a high precipice and A runs toward D threatening to tackle him in a way that would break several of D's ribs. If D shoots and kills A to prevent the tackle, this will count with more weight in the proportionality judgment than if D merely steps out of the way, allowing A to fall down the precipice and die.

14. *Whether the Aggressor Has a Duty of Care with Respect to the Person Who Will Bear the Threatened Harm*

It is sometimes thought that when D has a preexisting duty of care for the person who will suffer the threatened harm, this can provide D with agent-relative justification to inflict defensive harm on A that would not be possessed by someone who did not possess such duties of care. Whether or not this is correct, duties of care between D and the victim of threatened harm certainly affect whether D has a full liberty or a duty to engage in the defensive action.

V. LIABILITY AND LESSER EVIL JUSTIFICATIONS OF DEFENSIVE HARM

Fourteen complex factors affect liability to defensive harm. The first group of factors (1–8) concerns the evaluation of the agency of A; the second group of factors (9–14) concerns the evaluation of the agency of D. How should these factors be combined to ground conclusions about liability? What can we infer from those conclusions about lesser evil justifications?

I argued above that for A to be liable to defensive harm by D *just is* for that harm to be proportionate in the circumstances. On the analysis above, that suggests the following hypothesis: A is liable to defensive harm if and only if the combination of factors 1–8 exceeds, in the relevant way, the combination of factors 9–14.

There are, however, two reasons why this is not yet a sufficient characterization of the liability justification of defensive harm. First, more needs to be said about what is the “relevant way” in which the two groups of factors are to be compared. The second is that much the same characterization seems to be true also of lesser evil justifications of harm: D has a lesser evil justification for inflicting defensive harm on A if and only if factors 1–8 exceed in the relevant way factors 9–14.

How can that be? Lesser evil justification is an aspect of consequentialist ethics. Surely the only factors relevant to lesser evil justification are those concerning the outcomes of action: the magnitude and expected probability of the threatened and defensive harms, respectively. This may be true on a theory of strict act utilitarianism. But such a view is at odds with commonsense morality. For example, few people would think it is justified to inflict a lesser magnitude of harm on X in order to avert a greater magnitude of harm for Y, when the greater harm was a just punishment properly imposed on Y by X. Similarly, few would believe that it is justified on lesser evil grounds to kill a person who was about to cause by omission the death of several destitute people by cancelling a bank order to Oxfam. Whether action is justified on lesser evil grounds will in many cases depend on more than a simple comparison of magnitudes and probabilities of outcome harm and benefit.

Richer accounts of consequentialism, more commensurate with commonsense morality, will recognize that the value of states of affairs relevant to lesser evil justification depends crucially on how harms and benefits come into being, and how they are related to human agency. It is implicit in lesser evil justification that harm inflicted on an innocent bystander is a greater evil than an equivalent harm inflicted on a person liable to that harm. Similarly, harm brought about through positive agency is a greater evil than harm brought about through negative agency, and harm intentionally inflicted plausibly counts as a greater evil than harm brought about as a foreseen but unintended side effect of justified action. In fact, each of the factors that we have considered in the context of liability to defensive harm has relevance also to whether harm is justified as a lesser evil.

This is itself a striking conclusion. Despite their respective provenance in the divergent traditions of deontology and consequentialism, both the liability justification of defensive rights and the lesser evil justification have a similar underlying structure. They are linked by being in essence complex proportionality relationships between a shared set of relevant factors interpreted under differing conditions.

The vital question is: How do the combination of factors within these groups, and the comparison between groups, vary, to yield the undoubted differences between liability and lesser evil justifications? There are four key differences. First, different factors act as necessary and sufficient conditions of justification in liability and lesser evil accounts. Second, the two accounts assess benefits and harms to multiple agents differently. Third, they vary in their incorporation of benefits and harms to third parties. Fourth, factors can function as either thresholds or continuous contributors to justification, and do so dif-

ferently in liability assessments than when identifying lesser evils. I will explore each consideration in turn.

A. Necessary and Sufficient Conditions for Justification

Some of the fourteen factors are preconditions for the others playing any justificatory role. Preconditions in the first group (1–8) are necessary conditions for justification; preconditions in the second group (9–14) are factors whose absence is a sufficient condition for justification. Different factors are preconditions for liability and lesser evil justification.

The magnitude and probability of the threatened harm are both necessary conditions for lesser evil justification of harmful defensive action. Similarly, the magnitude of the defensive harm is a precondition in the second group: when D can avert a threatened harm without inflicting any harm at all, then it is necessarily a lesser evil and this is sufficient for justification. All other factors are additive: they increase or decrease the evil on either side of the relationship, but their being zero does not rule out or justify inflicting defensive harm.

For liability, two factors in the first group are necessary conditions of justification. First, A must have some responsibility for the threatened harm. Second, the threatened harm must be unjustified, on McMahan's view, or on my account a transgression of rights (including justified infringement). Other factors are additive, increasing or decreasing the weight of their side of the proportionality relation, but neither necessary nor sufficient for justification.

What about the magnitude and probability of the threatened harm? Surely it is a necessary condition of justification on the liability account that harmful defensive action responds to some nonzero value of threatened harm? I don't think that harm is necessary. D is justified in inflicting defensive harm if and when doing so averts a violation of rights (or, on my view, an infringement). D can be justified in inflicting some harm to prevent even a harmless rights violation—indeed, even one that actually benefits the right holder. Imagine D's financial representative can realize an immediate risk-free gain for D by making a certain trade. If D has instructed him not to do so, he is entitled to take action inflicting some harm on A to prevent him making the trade, even if his reason is sheer bloody-mindedness. D has a right that A not make that trade. Violating this right generates some liability to harm, though it is harmless, indeed beneficial to D.

B. Harms Inflicted on More than One Person

Liability and lesser evil justifications also diverge in the harms and benefits they consider relevant to justification. Liability is a localized comparison between persons in a situation of conflict; it concerns their interacting rights and duties and so values outside that relationship are irrelevant. Lesser evil justification has different constraints.

Two restrictions are particularly important. The first concerns how we treat defensive harms and threatened harms that affect more than one person. Consider defensive harm inflicted on multiple attackers. Within a liability justification, harms inflicted on multiple As are not aggregated, but considered separately. This is why inflicting defensive harm on *any* number of persons who are individually liable to that harm can be proportionate on a liability account.

Lesser evil justification, by contrast, aggregates the defensive harms inflicted on all affected persons. It discounts the evil attributed to harm inflicted on the liable, but unless the harm is discounted to zero, it is still possible that defensive harm inflicted on multiple liable persons will not be the lesser evil. This is one respect in which deontological reasoning is more permissive of the infliction of instrumental harm than consequentialist reasoning.

McMahan seems misled on this point. He says “harms to which people are liable do not count among the bad effects in any proportionality calculation.”²⁸ But that is not true if we are considering proportionality in lesser evil justification (wide proportionality). Indeed, McMahan himself recognizes this, earlier in the same chapter: “Harms to which people are liable are bad not only for those who suffer them but also from an impersonal point of view. Although their weight is discounted in proportionality calculations, they are never of merely neutral or positive value.”²⁹

Consider McMahan’s own example of ten innocent persons imprisoned by 500 soldiers. McMahan discusses why it seems impermissible to kill the 500 soldiers in order to free the prisoners, even though each soldier considered individually would be liable to be killed if necessary to free the prisoners. His solution attributes only a portion of the harm of wrongful imprisonment to each soldier, so that each of the 500 guards falls below the threshold for liability to be killed. But a more plausible analysis is that the 500 soldiers are indeed liable; however, killing them is all-things-considered impermissible for reasons of lesser evil. Even though we discount the evil attributed to killing a person who is liable to be killed, when the numbers are high enough, the aggregate evil can still provide com-

28. McMahan, *Killing in War*, 25.

29. *Ibid.*, 8.

pling reasons against inflicting the harm, potentially to the point of making self-sacrifice the obligatory lesser evil. The case is thus analogous to my analysis of civilians' right of defense against permissible collateral harm by a justified tactical bomber above.

Defensive harm inflicted on multiple aggressors is individuated for liability justification, but aggregated for lesser evil justification. What if the threatened harm is inflicted on multiple victims? Clearly this must be aggregated for lesser evil justification. But the same also seems true for liability. Suppose it would be disproportionate for D to kill A to prevent him breaking D's arm. If A's action would also break other people's arms, this is clearly relevant to whether A is liable to defensive harm. At some point, the cumulative harm of the broken arms may make a lethal response proportionate.

Arguably unjust threatened harm inflicted by A on multiple victims is also aggregated for determining liability even when not caused by a single act of A. Suppose A unjustifiably attacks D, and D can defend himself only by breaking A's arm. Suppose D also knows that A is an abusive husband who was at that moment returning home to beat his wife. That breaking A's arm will also prevent him from beating his wife is surely relevant to the proportionality of inflicting the defensive harm (provided that D was acting with the intention to defend both himself and the wife).

However, trivial threatened harms may be an exception to this conclusion. For example, if A were responsible for emitting a sound that irritates millions of people, if we aggregate all these harms, we might, counterintuitively, render A liable to suffer a severe defensive harm. Although I cannot settle the question here, it seems plausible that the aggregation of threatened harm across multiple victims becomes problematic only when it would change not simply the magnitude but the *kind* of defensive harm inflicted in response.

C. When Do Benefits "Offset" Harms for the Purposes of Proportionality?

Harmful actions sometimes also produce benefits. When can those benefits "offset" harms in the proportionality calculation? Here again there are differences between lesser evil and liability justification.

(i) *Offsetting benefits in liability justification.*—Liability to defensive harm is insensitive to benefits generated for third parties. Suppose A lethally attacks D, in circumstances where D's organs (but not A's) could save a sick patient who will otherwise certainly die. That D's death would benefit the patient does not mitigate A's liability to defensive harm from D. Benefits for third parties of A's harmful action cannot offset the threatened harm in liability justifications.

Similarly, that A's organs (but not D's) could save a patient's life does not augment his liability to defensive force if the threatened

harm he poses to D is insufficiently great. Benefits for third parties of D's harmful action also cannot offset defensive harm for proportionality calculations within liability justifications (with one important exception, noted below).

This has an interesting implication for whether one can become liable to defensive harm for all-things-considered justified rights infringements. A's justification for infringing D's right against harm often consists in the benefits that harming D has for third parties. Yet, as we have just seen, benefits that accrue to third parties from inflicting threatened harm are in general irrelevant to whether A is liable to defensive harm. This strongly suggests that A can be liable to defensive harm even when the threatened harm for which A is responsible is justified as the lesser evil.

Sometimes a threatened harm can generate benefits for A or D themselves. As with benefits to third parties, benefits that accrue to A or D are also generally irrelevant to A's liability to defensive harm. Quite obviously, the fact that A will benefit from harming D is irrelevant to his liability. The same is true for benefits to D from the threatened harm. If A unjustly tampers with D's car, causing \$100 in damage, that he prevents D from feeding his gambling addiction by attending a poker game at which he is almost certain to lose \$1,000 is irrelevant to A's liability to defensive force from D. However, intention and justification will play a role here. If A is D's addiction therapist, and he acts with the intention of stopping D from attending the game, then this can clearly be relevant to his liability by providing a potential justification for his action.

There is, however, an important exception to the general claim that the beneficial consequences of harmful action are irrelevant to liability justification. As we have seen, when action causes threatened harm for multiple victims, this is aggregated to determine liability. But when defensive action averts a threatened harm that would otherwise befall third parties, that is a benefit for them. Whether we describe this effect on third parties as augmenting the value of the threatened harm or as offsetting the value of the defensive harm for the purposes of proportionality seems irrelevant. The point is that there *is* a form of welfare effect for third parties that is relevant to A's liability.

The principle that determines when such further welfare effects are relevant to liability is clearly the following: benefits of D's harm-producing action offset defensive harm in determining A's liability if and only if those benefits consist precisely in averting the threatened harm for which A is responsible. Call this the *principle of offsetting benefits for liability*. We have seen this principle at work with benefits that accrue to third parties, but the principle applies also to benefits

of defensive harm that would flow to D himself. The only benefits to D that offset defensive harm are the benefits of averting the unjust threatened harm for which A is responsible.

What underlies the principle of offsetting benefits for liability? Again, this stems from the basic nature of the liability justification. A is liable to be harmfully used to avert only unjust harms for which he is responsible. It is therefore *his agency* that determines the relevant harms that fall within the scope of threatened harm for assessing his liability (or to express the same point in different terms, it is his agency that determines the relevant benefits that are capable of offsetting defensive harm for the purposes of assessing liability). In Section V.C.iii below, I explore the consequences of this principle for proportionality in just war theory.

(ii) *Offsetting benefits in lesser evil justification.*—There are substantial limits on how benefits offset harms for assessments of liability. But these same limits do not apply to lesser evil justification. Weighing the infliction of harm against the creation of benefits for third parties is a characteristic form of lesser evil reasoning. For example, in the organ transplant cases considered in the last section, the benefits for third parties clearly would be relevant to whether A or D could inflict harm as a lesser evil, even though they are irrelevant to liability. Indeed, one might think lesser evil justification should consider all consequences of action, and so must allow any benefit to offset harms for the purposes of proportionality. But this is not the case.

Consider a man who has promised his wife to be home by 6 p.m. He can get there on time only if he drives at 200 miles per hour through the center of town, knowing there is a significant risk that he will strike and kill a pedestrian. However, he is a surgeon at the local hospital, and he knows that there is a patient who will imminently die for want of an organ transplant. He reasons that if he kills a pedestrian, he will be able to use the cadaver's organs to save that patient. Even if this were true, he still acts impermissibly if he drives recklessly to make his rendezvous. But it is not easy to explain why. After all, the intended consequence of his action is good (even mildly obligatory), and if benefits offset harms, then the net unintended consequences of his action are morally neutral.

The answer must be that not all benefits offset harms in lesser evil justification. Consider a variant of the case. If the driver was racing home to pick up his surgeon wife to take her to the hospital to undertake a lifesaving operation, this benefit could offset the risk imposed on pedestrians. The relevant principle seems to be that when justifying unintended harms as the lesser evil, only the beneficial consequences of the intended objective of action can offset the unintended harms for proportionality. The beneficial consequences of the

unintended harm cannot offset the unintended harm. Call this the *principle of offsetting benefits for lesser evil*.

What underlies this principle? As we have already seen, intentions matter, even for lesser evil justification. Unintended harms are easier to justify because they do not instrumentalize their victims in the way that intended harms do. But the driver in the first variant of the case *does* instrumentalize the pedestrian by exposing him to harm. He has a conditional intention to use the pedestrian's death as a means to achieve benefits that play a necessary role in the justification of his action. That is why the benefits that follow from the harmful unintended consequences of his action are excluded.³⁰

(iii) *Sufficient just cause and contributing just cause in just war theory*.—The principles of offsetting benefits for liability and offsetting benefits for lesser evil can help to resolve a puzzle in just war theory. Several authors have noted that there are implicit restrictions on the benefits that can count toward proportionality judgments in war. Thomas Hurka, following an earlier discussion by McMahan and McKim, notes that only benefits that are contained within the sufficient or contributing just causes for war can contribute to proportionality.³¹ Sufficient just causes are those that suffice by themselves to satisfy the just cause condition, for example, resisting aggression or halting atrocity. Contributing just causes do not on their own suffice to satisfy the just cause condition, but once a sufficient just cause is present, they can contribute toward justification. Deterring future aggression, counteracting economic shocks caused by an aggressor's disruption of trade, and bringing to an end lesser forms of internal oppression that would not by itself justify humanitarian intervention are often thought to be contributing just causes. But certain other goods that seem qualitatively similar are excluded from counting as contributing just causes. For example, Hurka supposes that the coalition assembled to oust Iraqi forces from Kuwait in 1992 might also have had the effect of enabling a peace settlement in Israel. Yet intuitively, this good ef-

30. This argument is clearly in tension with Frances Kamm's work on the "doctrine of triple effect" (see *Intricate Ethics* [Oxford: Oxford University Press, 2007], chaps. 4 and 5). Kamm argues that it is possible to act *because* some harmful effect will occur, without *intending* those harms. However, this conclusion is consistent with the claim that acting because harmful effects will occur does objectionably instrumentalize the victim of these harms, even if it does so less than directly intended harm. Moreover, even if Kamm is correct that such conditional agency does not violate the prohibition on intentionally bringing about harm, the preceding argument suggests that the benefits that flow from such unintended harms are nonetheless excluded from offsetting harm in lesser evil justification.

31. Thomas Hurka, "Proportionality in the Morality of War," *Philosophy and Public Affairs* 33 (2004): 34–65; Jeff McMahan and Robert McKim, "The Just War and the Gulf War," *Canadian Journal of Philosophy* 23 (1993): 501–41.

fect of the war (supposing it had happened) could not offset the harms of war. Similarly, if going to war boosts economic production more generally, this good does not offset the harm of war.

But Hurka struggles to explain why this is so. He says: "This raises the question whether there is some unifying feature that gives these contributing causes their status. So far as I can see, there is not; like the sufficient just causes, they are just the items on a list."³² But these are not merely items on a list; what goes on the list is determined by the two principles we have identified. When establishing liability, the only benefits that can offset defensive harm are those that consist in averting unjust threatened harms for which A is responsible. That is why ending lesser forms of oppression and restoring the economic costs of aggression are relevant goods for establishing liability: they avert harms for which the aggressor is responsible. Deterring future aggression is a more difficult case, but if we plausibly assume that the harm of unjust aggression consists not only in the immediate threat posed, but also in weakening the norm of nonaggression, then clearly averting this threat should be part of the proportionality calculation. Conversely, stimulating economic production or facilitating a peace settlement in Israel are not goods that consist in averting harms for which the aggressor is responsible. That is why they are excluded from the proportionality calculation.

Lesser evil proportionality (wide proportionality) is subject to a different and more inclusive principle for determining which benefits can offset harms for the purposes of proportionality. Some benefits that do not consist merely in averting unjust threatened harm for which A is responsible will be relevant to wide proportionality in war. But there are still limits on the benefits that can offset harms for lesser evil justification. Unintended harm inflicted on A cannot be made proportionate by the beneficial side effects of that unintended harm. For example, in the 2008 Israeli campaign in Gaza, Operation Cast Lead, a number of Israeli officials suggested that although the collateral harms inflicted on Palestinian residents of Gaza during the operation were not directly intended by the Israel Defense Forces, they nonetheless contributed to the effectiveness of the operation by creating a disincentive for the population to support future Hamas attacks against Israel. But if the principle of offsetting benefits for lesser evil is correct, then such consequences cannot contribute to a lesser evil justification for the unintended harming of civilians. Such benefits are excluded from wide proportionality because they derive from the unintended harmful consequences of military action.

32. Hurka, "Proportionality in the Morality of War," 43.

D. Thresholds or Continuums?

Factors in the liability and lesser evil proportionality relationships will also vary depending on whether they admit of degrees or are a binary threshold. In liability justifications, some factors seem to be scalar in some cases, but are thresholds in others. Specifically, the magnitude and probability of defensive harm and of threatened harm and responsibility for threatened harm appear scalar when they involve damage to property and other fungible goods. For example, if action for which A is responsible threatens to wrongfully cost D \$1,000, A would be liable to a defensive harm of \$1,000. But if the probability of the harm befalling D absent defensive action is 50 percent, then A would be liable to \$500 worth of harm, likewise if A had some excuse or partially exculpating justification, or if half of the loss resulted from another party's intervening agency, or if D were responsible for half of the loss.

However, this picture is far less plausible when one considers harms to goods that are nonfungible, preeminently the loss of human life and grievous bodily injury. McMahan makes a heroic effort to make responsibility a scalar contributor to liability even for lethal threats.³³ However, there are obvious examples in which factors such as responsibility and probability are simple thresholds in the context of threats to life or grievous harm. For example, D has the right to use lethal force against the unjustified lethal attacks of a criminally insane psychopath, or an assassin acting under a fully exculpating excuse of duress.³⁴ Similarly, in the context of probability, D has the right to use lethal force against A if A was inflicting a small risk of death by forcing him to play Russian roulette. A remains fully liable to attack even though his responsibility, or the probability of the attack succeeding, may be barely more than zero. In other words, these factors are behaving, in the context of lethal harms, not as continuum contributors to liability but as thresholds for liability.

Why are factors like responsibility and probability scalar in some contexts and binary in others? The answer may depend on the nature of the threatened harm itself, in particular on whether it is possible to compensate D after the fact.

Liability to defensive harm is only one potential normative consequence of action that transgresses, or threatens to transgress, rights. The other potential normative consequences are liability to punitive harm, and liability to redress harm, most notably to provide compen-

33. See McMahan, *Killing in War*, chap. 4, esp. 192ff. McMahan is at his most persuasive in his sensitive discussion of child soldiers.

34. Assuming there is such a thing. Duress is not an excuse for homicide in many jurisdictions.

sation. These three forms of liability are linked by their shared origin in a relationship between the agency of A and D. Each seeks to rectify injury imposed by unjust action. Defensive rights and redress rights in particular are closely linked. Both provide a personal remedy for D to secure his rights against infringement (in this respect, both differ from liability to punitive harm, which is premised on the violation rather than merely the infringement of rights, and has as its primary goal the securing of rights on a broader, societal level). However, the two rights do this in quite different ways. Defensive rights are forward-looking and seek to preventively avert an unjust harm before it occurs. Redress rights are backward-looking and seek to repair or restore an unjust harm after it has occurred.

However, self-defense is a morally risky activity, for three principal reasons. First, we must act quickly, with little time for reflection. Second, our own interests are at stake, so we are unlikely to be properly impartial. Third, determining liability is a complex matter, and we frequently lack sufficient information to draw confident conclusions. When we inflict defensive harms, then, we run a substantial risk of wrongdoing. In a just society, by contrast, redress can remedy each of these problems: we have the time, the epistemic tools, and the impartiality to properly establish liability.

Removing these moral risks generates reasons of justice for A to defer vindication of his rights from defense to redress. It is plausible that these reasons increase in strength in precisely those cases where the risks are greatest, that is, in those cases in which the responsibility, or agency of A, or the probability of unjust harm is most uncertain.

Moreover, there are prudential as well as moral reasons for exchanging defensive rights for redress rights, since this enables agents to eliminate certain forms of personal risk. Defensive rights are liberties to undertake harmful acts that would normally be prohibited. Exercising a right of private defense therefore involves incurring a significant personal risk: if one makes an objectively unreasonable mistake as to the existence or extent of defensive rights, this can generate costly (sometimes criminal) liability. Deferring vindication of rights to the context of redress eliminates this risk.

Where harms are readily compensable, redress may be morally and prudentially preferable to defense in cases of uncertain liability. However, when there is no prospect of compensation for threatened harm, defense may be the morally preferred option. Although some think that all harms can in principle be compensated—by rendering the victim indifferent between {suffer harm, receive compensation} and {don't suffer harm}—this is at least false of death, and may well be false of some other harms too, because the victim can never be rendered indifferent in this way, either because other benefits are

incommensurable with the harm or because it is so severe that nothing can compensate. Death is the paradigm of a noncompensable harm because the wronged party no longer exists.

This helps explain why liability-determining factors like responsibility, justification, and probability of threatened harm are scalar in some contexts and binary in others. They are scalar when the threatened harm is a fungible good that can readily be compensated, but they function as thresholds for liability when the threatened harm is noncompensable, for example, the threat of death or other grievous bodily harm. The reason is that there are justice-based reasons to sometimes defer defensive rights for subsequent redress rights. But one cannot be required to defer vindication of rights to *post facto* redress when such redress is not possible, or is radically deficient, as it is in the case of harms to nonfungible goods.

VI. ALL-THINGS-CONSIDERED JUSTIFICATION

How do liability and lesser evil justifications interact to yield all-things-considered judgments of permissibility? In particular, when do lesser evil considerations provide an all-things-considered justification for infringing rights?

In general, we are much more inclined to allow lesser evil justifications for intentionally infringing property rights than for infringing the right to life or bodily integrity. This partly reflects the far greater moral importance of those rights compared with property rights. But this is not the whole story. It may be justifiable to intentionally burn an acre of someone's land to prevent a wildfire from engulfing five acres, but it would not be justifiable to intentionally kill an innocent person to save five others. In fact, once one puts fictional cases to one side, it is extremely difficult to find any example of all-things-considered justification for intentionally killing or grievously harming a person who is not liable to harm and who would not otherwise die or suffer the harm. In the most celebrated legal case to consider the lesser evil justification for homicide, *R v. Dudley and Stephens*, the two shipwrecked sailors were denied a justification for killing and eating the cabin boy, though it was reasonable to suppose they would all have otherwise have died.³⁵

The distinction between compensable and noncompensable harms helps to explain this difference. When a right is justifiably infringed as a lesser evil, it grounds a claim for compensation or redress. Infringing a right, even with the intention of fully compensating the right-bearer, still requires significant justification, because it imposes

35. *R v. Dudley and Stephens* [1884] 14 QBD 273 DC.

significant costs and risks on the right-holder—costs and risks he has a right not to bear. Moreover, unlike with deferred defensive rights, these costs are not balanced by any commensurate benefit of mitigating moral or legal risk for the right-holder. In lesser evil justification, it is not the right-holder, but the right-infringer, who assumes a risk of criminal liability if his interpretation of the justifying conditions proves to be defective. Nonetheless, with fungible goods the justified infringement of rights can be comprehended within the normative logic of rights because the infringed right does not simply disappear: it is transmuted into a right to compensation.

But in the case of noncompensable harms—grievous bodily injury, torture, rape, and, paradigmatically, death—this accommodation is not possible. Here the harm inflicted departs the system of rights altogether. The right has not simply been infringed to resurface as a ground for claims to restore the status quo ante; from the right-holder's perspective, the right (or a significant part of it) has disappeared. It is not surprising that intentionally transgressing such rights as a means to avert a greater evil represents an extraordinary taboo, never or almost never permitted in civilized systems of morality.

But this argument invites an important response. It is equally true of foreseen but unintended noncompensable harms, like the collateral killing of civilians in war, that the harm inflicted has departed the system of rights altogether. Why should it sometimes be all-things-considered permissible to unintentionally inflict noncompensable harm as the lesser evil?³⁶

I believe that we have indeed historically been overpermissive of the unintended killing of civilians in war. The prohibition on collateral killing should be much closer to the prohibition on intentional killing, than current norms recognize. Intention can make a difference to permissibility, and there clearly are circumstances in which it is all-things-considered justifiable to act in a way that foreseeably leads to the death of nonliable persons. Consider a road planner who foresees that a number of persons will die in collisions if a road is built who would otherwise remain alive. The lesser evil justification in this case is supported by a multitude of interrelated mitigating factors beyond the unintended nature of the harm. First, the harm is causally remote (it is not the construction of the road that kills the drivers). Second, the harm derives from the intervening agency of individual drivers who voluntarily take to the road knowing of the risks. Third, the drivers who bear the risks of the road also reap its benefits. Fourth, in a democratic society, drivers have an opportunity to par-

36. I am grateful to a reviewer for this objection.

participate in shaping the risk/reward trade-off implicit in the design of a road, thus imparting an element of consent to the risks they bear.

Contrast this with a typical case of collateral killing in war. Beyond the magnitude of the goods at stake (which may indeed be great), the single mitigating factor is the narrowly unintended nature of the harm. The fact that the harm is not directly intended does create at least the possibility of justification. But the conditions for permissibly infringing noncompensable rights remain extraordinarily restrictive. Consider comparable situations in domestic society. In a case of domestic self-defense or protective police action, it would not be permissible to unintentionally shoot one nonliable person as a side effect of saving another. In standard formulations of the Trolley Problem, it is assumed that saving five lives is required to justify the unintentional killing of one nonliable person. The special status of noncompensable harm partially explains this highly restrictive standard. Persons whose rights are infringed are owed compensation; if compensation is not possible, this provides compelling additional reasons against transgressing the right.

Why, then, are commonly accepted standards for the permissibility of collateral damage in war so much more lenient? A plausible hypothesis is that war planners have historically discounted the rights and interests of enemy noncombatants, because they have implicitly attributed partial responsibility to them for the harms that constitute the just cause (a thought that gains some support from McMahan's second form of narrow proportionality). Less charitably, their rights may be discounted simply because they are foreigners. Although I cannot argue the case here, I believe that neither rationale withstands rigorous assessment and that we must instead very substantially tighten restrictions on unintentional killing in war so as to bring them more into line with domestic self-defense and rescue norms.

VII. CONCLUSIONS

McMahan's identification of wide and narrow proportionality constraints lays the foundations for a more nuanced and sophisticated understanding of justifications for inflicting harm. Developing this central insight can lead to a valuable and surprising conclusion: that liability and lesser evil justifications in fact share much in common, in essence being complex forms of proportionality relationship between a shared set of underlying factors.

What, though, are the implications of this analysis for McMahan's justifications for killing in war specifically? Certainly the critique of the "moral equality of combatants" seems generally sound: most combatants on the unjust side pose rights-violating threats to their adver-

saries; if the latter are justified in using defensive force, the threats they pose are neither rights-violating nor rights-infringing, so unjust combatants cannot appeal to liability-based justifications for killing just combatants. Nor will lesser evil justifications be readily available, when the harms inflicted are noncompensable (and egregious). Besides undermining unjust combatants' rights to fight, McMahan has also argued for granting extended permissions to just combatants—specifically, allowing them in some cases to target noncombatants. The foregoing analysis of the importance of causal proximity to liability determinations casts some doubt on this view. Any contributions made by noncombatants to threats posed by their state will be far down the causal chain and may not be sufficient to ground liability. Taken together, these results support a view that I have elsewhere called “restrictive asymmetry”—denying that unjust combatants possess full combatant rights while affirming the protected status of all noncombatants.³⁷

Ultimately, however, two larger questions loom. The first is whether either the liability or the lesser evil paradigms can plausibly ground the justification of killing in war. I have written skeptically on that question before.³⁸ The sheer complexity of considerations that determine liability and lesser evil justification, and the numerous ways that defensive rights are constrained and limited, may mean that the ultimate implication of this analysis is to endorse a form of contingent pacifism, or at the very least to suggest far-reaching modifications in the way we ought to plan for, equip, and conduct military operations.

Second, this analysis has shown consequentialist lesser evil reasoning and deontological liability reasoning to be closer and more deeply interpenetrated than is typically believed. How far can this process of integration be taken? Are remaining conflicts between the two forms of justification merely evidence that they have not been suitably qualified and developed? Is a “grand unified theory” possible? This analysis points tantalizingly in this direction. But much further work will be required before these questions can be definitively answered.

37. Rodin, “The Moral Inequality of Soldiers.”

38. Rodin, *War and Self-Defense*.