

1 LIABILITY, PROPORTIONALITY, AND THE NUMBER OF AGGRESSORS

Jeff McMahan

I. Liability Justifications and Numbers

Proportionality can be a demanding constraint on the use of force. It can prohibit an innocent person from engaging in necessary defensive action. Suppose that the only means by which a person can prevent herself from being killed by a culpable attacker would unavoidably kill five innocent bystanders as a side effect. This would be disproportionate, so the victim must submit to being wrongly killed. Defensive action may also be disproportionate because of its effect on the threatener. If the only way one can prevent oneself from being wrongly and painfully pinched is to kill the pincher, one must submit to being pinched.

The form of proportionality that weighs an act's good effects against the harms it causes to victims who are not liable to those harms is *wide* proportionality. Everyone agrees that wide proportionality is relevant to the morality of war, so that an act of war can be impermissible because the harm it would inflict on civilian bystanders would be excessive in relation to the act's good effects. Few people, however, accept that a war or act of war could be disproportionate just because of the harm it would cause to unjust aggressors. Yet, given that outside the context of war there can be disproportionality in harms caused to threateners, such as the culpable pincher, who are potentially liable to *some* degree of harm, there is no reason why the same cannot be true in war. Proportionality in harms to those who are liable to some degree of harm is *narrow* proportionality. (As I will indicate in section V, it can be permissible to inflict harm that is disproportionate in the narrow sense—that is, in excess of the harm to which the victim is liable—if that harm is nonetheless proportionate in the wide sense.)

In practice, disproportionality in war is usually a matter of the *number* of people harmed or killed. If an act of war would be disproportionate in the

wide sense, that is generally because it would kill or injure too many civilian bystanders. And one might similarly suppose that a war or act of war could be disproportionate in the narrow sense because it would kill too many combatants (or, rather, too many combatants on the unjust side, on the assumption that combatants on the just side are seldom liable to attack at all). One might think, for example, that the Falklands War either was narrowly disproportionate or would have been if it had been necessary to kill a much larger number of Argentine combatants to preserve British sovereignty over the islands, as there must be some limit to the number of such people it could have been permissible to kill in pursuit of that arguably not very significant goal.

Although I have been attracted to this understanding of narrow proportionality, I now think it is mistaken. Narrow proportionality is a constraint on individual liability. It sets the limit to the amount of harm it can be permissible to inflict on an individual on grounds of liability. It has no application to numbers of individuals, unless a collection of individuals itself constitutes an individual that can be liable—an idea that I reject but cannot discuss here. Narrow proportionality in war is thus not sensitive to the number of combatants killed.

Yet, as the example of a hypothetical Falklands War in which it would be necessary to kill a vast number of Argentine combatants suggests, there must, at least in many cases, be some limit to the number even of unjust combatants that it can be permissible to kill in pursuit of a just cause. I will argue in the final section that to accommodate this fact we must recognize a further distinct form of proportionality, in addition to the wide and narrow forms.

(Some might object that if there were a proportionality limit to the killing of combatants, an unjust aggressor could then assemble sufficiently many conscripts, perhaps including child soldiers, to make defensive resistance disproportionate, thereby inducing moral paralysis in a scrupulous adversary. I think, however, that this is not an objection but a fact that we must accept, just as we must accept that an aggressor can make defensive action disproportionate in the wide sense by using a sufficient number of innocent shields.)

It will be helpful to examine four simplified examples rather than continuing to discuss the killing of combatants in war. This may prevent us from being distracted by intuitions specifically concerned with war that may reflect common but mistaken views about the ethics of war.

Suppose that a morally innocent and unthreatening person will be unjustifiably killed by a fully culpable aggressor unless the aggressor is killed. It seems permissible for the victim or a third party to kill the culpable aggressor, if this is the only way to prevent him from killing the victim, and if there would be no bad side effects that would outweigh the saving of the victim's life.

The justification for the defensive killing of the culpable aggressor is not that he deserves to die, or that killing him would have better consequences, or that

killing him would be the lesser evil, in that the only alternative would involve substantially more harm to others. It is not that he has consented to be killed, or that the victim is permitted to give priority to her own life over his. Nor is it simply that he is the aggressor. It is instead that, by being the responsible, *unjustified* aggressor, he has made himself *morally liable* to be killed. For the victim or a third party, it is unavoidable that either the victim or the aggressor will be killed. The aggressor is responsible for this fact while the victim is not; it is, therefore, a matter of justice that the aggressor should incur the costs of his own voluntary, unjustified action. Provided there is no other compelling reason why the unavoidable harm should go to the victim instead, the victim or a third party is justified in killing the aggressor, who has forfeited his right not to be killed in the circumstances.

Next imagine that a thousand such culpable killers will attempt to kill the same innocent victim, each appearing one after the other, in rapid succession. Suppose the victim knows that if she kills the first one, she will immediately be killed by the second unless she kills him as well, and so on. Suppose further that there would be no interval of life worth living between the killings. Killing one killer would enable her to have only a second or two of terror before the next arrived. Only if she kills all thousand will she have any further life worth living. Finally, suppose that none of the culpable killers is specially related to anyone who would be harmed as a side effect of his being killed, and that none of them would again pose a threat to others.

I suspect that most people would accept that it is permissible for the victim to kill all thousand culpable potential killers (hereafter “culpable killers”). It may seem that an innocent and unthreatening person is never morally required to submit to being wrongly killed just to spare one or more people who would otherwise culpably kill her. In each such case, the culpable killer seems to make himself liable to be killed.

Liability justifications appear to work by “pairwise comparisons” – that is, by considering each choice between harmings on its own and determining on which of the parties the unavoidable harm ought, as a matter of justice, to be imposed. (This is part of the explanation of why narrow proportionality does not take account of numbers.) In each choice, it seems that the culpable killer is liable to be killed whereas the victim is not.

Next imagine a case like the first except that the aggressor is not culpable. Although he is responsible for the threat he poses, he is only minimally responsible, by which I mean “merely responsible,” or responsible though not culpable. He might, for example, have chosen to engage in a permissible type of action, such as driving a car, that involves a very tiny risk of inflicting great harm on one or more innocent people. Through sheer bad luck, he now threatens the life of an innocent bystander.

The potential killer in this third example is a “minimally responsible killer” (hereafter, for brevity, “responsible killer”). Assume that he and the potential victim would suffer a roughly equal loss in being killed. Many people believe that it is morally permissible for the potential victim or a third party to kill the responsible killer if that is necessary to prevent him from killing the victim. This is because the responsible killer is liable to be killed in the circumstances. He has voluntarily chosen to engage in a morally optional activity that involves a reasonably foreseeable risk of causing great harm to others. If through bad luck his choice results in a threat to another’s life, the fact that he, but not the victim, bears some responsibility for the fact that one of them must die provides a reason of justice for imposing the harm on him. Assuming that the harm they would suffer in being killed is much the same, and that any side effects would also be roughly equal, there is a liability justification for killing the responsible killer.¹

The fourth and final example is like the second except that each of the thousand killers is only minimally responsible rather than culpable. Assume that each potential killer would suffer the same loss in dying as the victim and that the killers do not act in coordination but independently. If there is a liability justification for killing a single responsible killer, and if there is a liability justification for killing each of the thousand culpable killers, then by parity of reasoning there should also be a liability justification for killing each of the thousand responsible killers. Yet I find this intuitively implausible, and it would remain implausible even if the number of responsible killers was significantly lower—for example, a hundred. I have chosen the higher figure of a thousand precisely to try to elicit the intuition that killing that many responsible killers would be impermissible.

Even more troublingly, if there is a liability justification in each pairwise comparison that is unaffected by the number of killers, there is then *no* limit to the number of killers it would be permissible to kill in defense of a single victim. Many people may find this acceptable in the case of culpable killers, but it is implausible in the case of responsible killers, as the relevant difference between each responsible killer and the innocent victim is small—so small that some philosophers believe that even a single responsible killer on his own is not liable to be killed.

If one thinks that there is a limit to the number of responsible killers it can be permissible to kill in defense of a single victim, but no limit, or at least a higher limit, to the number of culpable killers it can be permissible to kill, one must explain how this can be so, particularly if the justification for killing potential killers of each type is the same – namely, a liability justification. More generally, if one

1. Jeff McMahan, *Killing in War* (Oxford: Clarendon Press, 2009), 166.

accepts that there is a liability justification for killing a single culpable killer and for killing a single responsible killer, and if both justifications work via pairwise comparisons, one must explain how there could be a limit to the number of responsible killers it can be permissible to kill. For if liability justifications work via individual pairwise comparisons, it seems that the justification for killing one aggressor should in general be independent of the justification for killing another, so that the number of aggressors ought not to matter to whether there is a liability justification for killing each. It may seem that whether one potential killer is liable to be killed depends only on facts about his own action and not on what other people might do independently. Furthermore, if a person's being liable to be killed is a defeasible justification for killing him, and if each responsible killer is liable to be killed, it should be permissible to kill all thousand, in the absence of countervailing considerations.

One might, of course, simply accept that there is no limit to the number of responsible killers that it can be permissible to kill in defense of a single innocent victim. One might argue that rights and liabilities are moral facts about relations between individuals, so that if one individual threatens to violate another's rights, he becomes morally liable to necessary and proportionate defensive force irrespective of how many others may threaten the same rights of the same victim.² If, moreover, the right that is threatened is the right not to be killed, killing the threatener should be proportionate.

Several prominent philosophers seem to accept this view. Frances Kamm writes that

a response to multiple wrongdoers can satisfy narrow proportionality so long as the response of each is proportional to his wrongdoing. This is very clear in a domestic case, for if each of many people is trying to paralyze you... it could be a proportionate response to kill all the wrongdoers to prevent the paralysis of one person. On the basis of this sort of case, one might describe the determination of a proportional response to wrongdoing as involving "pairwise comparison"...: One compares the wrong to be avoided with what would have to be done to each wrongdoer one at a time, and if there is no violation of proportionality in any individual comparison then there is no violation tout court.³

2. David Rodin, *War and Self-Defense* (Oxford: Clarendon Press, 2002).

3. F.M. Kamm, *Ethics for Enemies: Terror, Torture, and War* (Oxford: Oxford University Press, 2011), 133–34.

By stipulating that the many people are *trying* to cause the paralysis, Kamm presumably intends to suggest that they are all culpable. But her point about pairwise comparisons applies equally to merely responsible killers, as does her assumption that proportionality is not a relation between harms caused and harm prevented but between harms caused and the *wrong* “to be avoided.”⁴ So the implication still seems to be that if killing each responsible killer is proportionate, so that each is liable to be killed independently of what the others may do, there is no limit to the number it can be proportionate and permissible to kill in defense of the single victim.

II. Effectiveness as a Condition of Liability to Harm

If we assume, as I think we should, that proportionality weighs the harm an act will cause against the act’s good effects, particularly the harms it will prevent, rather than against the strength of the rights the act might protect, then there is a simple way to avoid the conclusion that there is no limit to the number of responsible killers it can be permissible to kill. I have stipulated that the tiny interval of life the victim would secure by killing any one responsible killer before either killing or being killed by the next would not be worth living. Her killing any one responsible killer would thus be defensively ineffective. On its own it would do no good. As I understand liability, it is unlike desert because it is essentially *instrumental*, in that a person cannot be liable to be harmed if harming him will be neither a means nor a side effect of achieving some good effect. Effectiveness is thus a necessary condition of liability and if killing a responsible killer would be wholly ineffective, he is not liable to be killed.

There are ways in which killing a responsible killer might be effective other than as physical defense. Killing him might be an effective assertion of the victim’s moral status, or it might preserve deterrence. Or we could suppose that each killing would afford the victim some short interval of life worth living. But even if one or all of these conditions obtained, the killing would still be narrowly disproportionate. This is in part because the killer’s minimal responsibility diminishes the offense against the victim’s moral status implied by his action and in part because efforts to preserve deterrence are less effective the less responsible potential threateners are. But if killing the responsible killer would be narrowly disproportionate, that too means that the killer is not liable to be killed; for, as I understand liability, it is a necessary truth that one cannot be liable to a harm that is

4. For a challenge to this latter claim, see Jeff McMahan, “The Limits of Self-Defense,” in *The Ethics of Self-Defense*, ed. Christian Coons and Michael Weber (New York: Oxford University Press, 2016), 208–10.

disproportionate in the narrow sense—that is, a harm that exceeds that to which one is liable. Narrow proportionality, like effectiveness and necessity, is thus internal to liability.⁵ (There are, of course, conceptions of liability according to which effectiveness, necessity, and narrow proportionality are not internal but are constraints on the permissibility of acting on a liability justification. In section III.3, I will indicate one implication of this view.) Effectiveness, necessity, and narrow proportionality are what may be called the “circumstantial conditions” of liability, in contrast with the “agential conditions,” which for present purposes we may take to be causal and moral responsibility for a threat of wrongful harm.

It may seem that it cannot be right that none of the responsible killers is liable to be killed. For the same logic applies to killing the thousand culpable killers. If killing a single responsible killer would be ineffective, so would killing a single culpable killer. And I suspect that most people believe that a single culpable killer would not be wronged by being killed, even if the victim would then be immediately killed by a different culpable killer. I believe, by contrast, that even a culpable killer has a right not to be killed gratuitously, without producing any good effect, as would be the case if he alone were killed in a case involving overdetermination, such as the second case I introduced in section I. But I will not defend that claim here.⁶

I have thus far been assuming that each responsible killer, other than the last, is not liable to be killed because killing him would be defensively ineffective, given that the victim’s loss of further good life is overdetermined. This, however, leaves open the possibility that killing any individual responsible killer *would* be effective on condition that the other 999 killers are killed as well. But this is a problematic condition for the liability of each; for, assuming compliance with morality, all the others *will* be killed only if it is *permissible* to kill them. Yet whether it is permissible to kill them seems to depend on whether each is liable to be killed. This is because a liability justification is the only possible form of justification for killing a thousand people as a means of saving only one. There is certainly no lesser-evil justification; neither do the responsible killers consent to be killed; nor does the victim have an agent-relative permission to give her life priority over those of a thousand people who are only minimally responsible for the threats they pose; and so on. If this is right, it can be permissible to kill all thousand only if each is morally liable to be killed; yet whether each is liable depends on whether it is permissible to kill them all. This seems viciously circular. I will return to this problem in section VI.

5. Ibid.

6. I have defended it in *ibid.*, 206–10.

III. Beneficence Overrides Liability?

David Rodin has offered an appealing explanation of how there can be a limit to the number of responsible killers it can be permissible to kill. He suggests that, even though each killer is morally liable to be killed, the many individual liability justifications are overridden by a reason of beneficence (or, perhaps more precisely, non-maleficence) not to cause so much harm as would be involved in killing all thousand responsible killers. He begins by endorsing the view that liability justifications work via pairwise comparisons, which seem unaffected by how many such justifications there may be.

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.... Within a liability justification, harms inflicted on multiple [aggressors] 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.⁷

Rodin does not discuss the problem of overdetermination but goes on to contrast liability justifications with a different form of justification—a “lesser-evil justification,” which

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.⁸

I interpret this to mean that, even though all of the responsible killers are liable to be killed, and even if their interests are accordingly discounted for their responsibility, killing all of them would be so much the greater evil that it would not be permissible. There is thus a lesser evil justification for allowing the victim to be killed rather than causing the vast harm involved in killing all thousand responsible killers. While normally lesser evil justifications override the constraint against actively harming or killing an innocent person, here there is a lesser evil justification for *allowing* a nonliable person to be killed when the alternative is *killing* a

7. David Rodin, “Justifying Harm,” *Ethics* 122 (2011): 74–110, at 99.

8. *Ibid.*

large number of people who are *liable* to be killed. The liability justifications are, as I have put it, overridden by a “greater-evil constraint.”⁹

In an earlier draft of this paper, I presented three objections to Rodin’s view. In his contribution to this book, Rodin has replied to these objections. I have therefore left the objections largely as they were in the draft to which Rodin has responded. I have also included some brief comments on his replies. I have, however, subsequently come to believe that his view and the view I will defend in section VI are quite similar. Our main differences concern, first, whether each responsible killer is liable on his own and thus whether there is a liability justification for killing each that must be overridden, and, second, whether the explanation of the impermissibility of killing them all also applies to the killing of all the culpable killers. I will consider the second in the following subsection and return to the first at the end.

III.1 Lesser Evil

One reason for skepticism about Rodin’s view is that it seems to apply equally to the killing of the culpable killers. The problem is not that, unless harms to the culpable killers are discounted to zero, there must be some limit to the number of killers it can be permissible to kill. That may well be true. The problem is instead that on Rodin’s view the responsibility (or culpability) of the killers has already been taken into account in the liability justifications. What is now supposed to weigh against those justifications is the combined harm that will be suffered by the many killers if all are killed. That is, considerations of well-being, or beneficence, now weigh against considerations of liability. But when the concern is with well-being rather than the bases of liability, the well-being of the culpable killers matters or counts in the same way as that of the responsible killers. While the culpability or responsibility of the killers is relevant to the determination of their liability in the establishment of a liability justification, it is not relevant to how much weight their well-being has in a lesser evil justification or greater evil constraint. Lesser evil justifications do not take account of responsibility or culpability precisely because they apply to harms to which the victims are not liable. It is for this reason that, if the individual liability justifications for killing the responsible killers are overridden by what Rodin now calls a “lesser evil obligation,” the same must be true of the liability justifications for killing the culpable killers.

Rodin writes in response to this objection that in a passage he quotes (identified by his note 11), I myself acknowledged that, in his words, “harm inflicted on a person who is liable to that harm is not as bad from an impersonal perspective

9. For further discussion, see Jeff McMahan, “Duty, Obedience, Desert, and Proportionality in War: A Response,” *Ethics* 122 (2011): 135–67, at 152–53.

as harm inflicted on a non-liaible person.” He then writes that “it therefore has less negative weight in lesser-evil reasoning.”¹⁰ This is not quite what I meant, though I concede that I did not express my meaning well. First, my reference to impersonal badness was unnecessary, as that was a side issue. In general, lesser evil justifications weigh harms against other harms. Thus, what seems to weigh against the killing of the responsible killers is primarily the harms they will suffer in being killed, not the impersonal badness of their absence from the world. There might, moreover, be impersonal considerations that favor killing them—for example, to paraphrase an old German legal maxim, that Right should not yield to Wrong.

Second, my claim in the passage he quotes is that a threatener’s interests are discounted relative to those of his victim in the assessment of proportionality *in a liability justification*. What I meant is that it is often permissible to inflict harm on a threatener that is considerably greater than the harm that the threatener would otherwise have inflicted on his victim. It is thus in the determination of how much harm it can be justifiable to inflict on a person *on grounds of liability* that considerations of responsibility and culpability are relevant. A threatener’s responsibility or culpability is one determinant of the degree of harm to which he is liable. It has then exhausted its power to justify the infliction of harm on the threatener. If we have determined the amount of harm to which a culpable threatener is liable but then come to believe that his culpability justifies harming him to an even greater degree, we should revise our judgment about the amount of harm to which he is liable.

III.2 The Justificatory Priority of Liability

A second ground of skepticism is that it is odd to suppose that it could be justifiable to inflict harms, or allow harms to be inflicted, on nonliable people as an alternative to inflicting harms on those who are liable to them. With two possible exceptions, a liability justification for the infliction of harm seems to exclude the permissibility of other distributions of the harm. This is because liability is a matter of justice in the distribution of unavoidable harm. Suppose that there are only two options for preventing a catastrophe. One is to kill some large number of people, all of whom are culpably responsible for the impending catastrophe and are thus liable to be killed as a means of preventing it. The other is to kill an innocent bystander as a means or side effect of preventing the action of those who are liable. It seems impermissible to sacrifice the one to spare the others from harms to which they have made themselves liable, even though this would be much the lesser evil.¹¹

10. David Rodin, “The Lesser Evil Obligation,” in this volume, _.

11. “The Limits of Self-Defense,” Section 1.

The two exceptions to which I referred are these. First, in some cases it can be permissible (though not necessarily required) to harm a nonliable person rather than a liable person if the former freely consents to suffer the harm. Second, a liability justification for harming a person can be overridden when harming him would cause disproportionate harm to innocent bystanders as a side effect.

Rodin has two responses to this objection. The first appeals to the example of a pauper who recklessly damages a millionaire's car. Rodin says that, although the pauper is liable to compensate the millionaire, the latter has a lesser evil obligation to waive his claim because compensation would provide only a marginal benefit to him but would involve a "devastating loss" to the pauper. This, he says, shows that considerations of "distributive justice" can be overridden by a lesser evil obligation. I think, by contrast, that this is simply a case in which it would be narrowly disproportionate to coerce the pauper to pay full compensation, even taking the pauper's recklessness into account. If the harm to the pauper would be disproportionate in relation to the benefit to the millionaire, and if narrow proportionality is internal to liability, then the pauper is not morally liable to provide full compensation (though he may be legally liable to do so).

Rodin's second response is to present an example that challenges my claim that liability is a matter of justice in the distribution of unavoidable harm. In the example, X wrongly provokes Y but Y reacts disproportionately by trying to kill X. X can save his life only by breaking Y's wrist, which would be narrowly proportionate (in the harm it would inflict on Y), but would also prevent Y from performing a life-saving surgery on Z that only Y can perform. Rodin comments that "the incidental harm inflicted on Z is proportionate in the wide sense since a) the unjust harm inflicted on Z is not greater than the unjust harm averted from X, and b) the harm inflicted on Z is intended by X neither as a means nor as an end in itself."¹²

If the harm inflicted on Z were proportionate in the wide sense, that would not be for the reasons Rodin cites. Those reasons would also apply if X were to *kill* Z as a side effect of defending himself against Y, but that, I believe, would be disproportionate in the wide sense. Wide proportionality is arguably sensitive to the difference between killing an innocent bystander and preventing a bystander from being saved.¹³

It is not, however, widely proportionate to prevent Z from being saved as a side effect of saving X. It is arguable, though controversial, that it is widely disproportionate to prevent an innocent person from being saved as a side effect of saving

12. Rodin, _.

13. Jeff McMahan, "Proportionate Defense," *Journal of Transnational Law and Policy* 23 (2013–14): 1–36, 34–35.

only one other innocent person, just as it is disproportionate to kill an innocent person as a side effect of saving only one innocent person. Wide proportionality is a constraint on a lesser evil justification for the infliction of harms to which the victims are not liable. When the deaths of two innocent people are equally harmful, neither seems to qualify as the lesser evil, no matter how it is brought about.

But what is decisive here is that X, but not Z, bears some responsibility for the fact that either he or Z must die, for he culpably provoked Y. Rodin in fact concedes that “if justice requires that unavoidable harm be borne by the party most responsible for the fact that the harm must be distributed, then as a matter of justice X rather than Z should bear the unavoidable harm.”¹⁴ Yet he thinks that this is incompatible with what my view implies about this case. He writes that if “liability is a matter of justice in the distribution of unavoidable harms, then it would not be the case that Y is liable to have his wrist broken” and “McMahan would appear to be committed to the conclusion that X is liable to be killed.”¹⁵

But my view implies neither of these claims. Assuming that Y is more responsible for the threat he poses to X’s life than X is on account of his provocation, my view implies that the breaking of Y’s wrist is a narrowly proportionate defense against a wrongful killing and thus that Y is liable to have his wrist broken. Yet, while this is so, X may not act on that liability justification because doing so would be disproportionate in the wide sense because it would prevent Z from being saved. Furthermore, X is not liable to be killed by Y or anyone else. But because X bears some responsibility for the fact that either he or Z must die while Z bears none, he *is* liable to be *allowed to be wrongly killed*, either by himself or by third parties, when the only options are saving him, thereby preventing Z from being saved, and allowing Z to be saved.

III.3 Liability as a Robust Justification

A third ground of skepticism about Rodin’s proposal is the most serious. As I noted, liability justifications can be overridden. They are only pro tanto justifications. They justify only the harm inflicted on the liable person. But they do unconditionally justify *that*. The person who is liable to be harmed has forfeited his right not to be harmed and there is a reason of justice to harm him, which is that if he is not harmed, harm will unavoidably befall someone else who bears no responsibility, or less responsibility, for the fact that someone must be harmed. Rodin’s claim, however, is that the harms to the many aggressors that have been justified on the ground that those aggressors are liable to them somehow also override those same

14. Rodin, _.

15. Ibid., _.

justifications. That is, the same harms for which there are liability justifications somehow defeat their own justifications. This seems incoherent.

Rodin's response to this is that "a justification is simply a consideration that defeats a moral objection." A liability justification, in particular, defeats the objection that a person has a right not to be harmed. Liability is the loss of such a right.

This seems right if effectiveness, necessity, and narrow proportionality are external constraints on a liability justification rather than circumstantial conditions of liability itself. That Rodin accepts this externalist understanding of what I take to be the circumstantial conditions of liability is suggested by his assumption that each responsible killer and each culpable killer is liable to be killed despite the overdetermination, which means that killing any one of them or any number fewer than all would be ineffective.

According to this understanding of liability, when a person becomes liable to be harmed, he has forfeited his right not to be harmed *in any way* or *to any degree*, even if harming him will have no good effects. Lacking a right, he cannot be wronged by the infliction of a harm that is narrowly disproportionate—even vastly disproportionate—and therefore wrong. For narrow proportionality is on this view an external constraint on the permissibility of acting on a liability justification. Nor can he be wronged by being harmed wholly gratuitously, for effectiveness is also only an external constraint.

For this and other reasons I have stated elsewhere, I think effectiveness, necessity, and narrow proportionality are internal to liability.¹⁶ According to this conception of liability, a liability justification is, as I claimed above, more than merely the elimination of one barrier to permissibility. For acting on such a justification is necessary, when some harm is unavoidable within a set of individuals, for the harm to go to the only individual or individuals who will not be wronged by being harmed. For this reason, an act for which there is a liability justification is generally morally required, unless it is widely disproportionate or excessively costly to the agent. And a moral requirement to do some act involves more than merely the defeat of a moral objection to that act.

IV. The Harm of Making People Liable to Be Harmed

An alternative and highly ingenious explanation of why it is impermissible to kill all thousand responsible killers has been defended by Kerah Gordon-Solmon.¹⁷ Gordon-Solmon accepts that, because of overdetermination, none of the killers is

16. "The Limits of Self-Defense."

17. Kerah Gordon-Solmon, "Self-Defense Against Multiple Threats," *Journal of Moral Philosophy* (published online 2015, awaiting publication).

liable to be killed in advance of defensive action. (Whether they are liable is, she says, indeterminate.) She concedes, however, that killing all of them would secure the conditions of their individual liability, as the killing of each would, as I have indicated, be effective, necessary, and narrowly proportionate *if* the other 999 were killed as well.

She then claims that it would be bad for the responsible killers to have “the liability justifications on the basis of which they’re killed secured.”¹⁸ Yet there is no right not to be made liable to be killed and the responsible killers cannot forfeit a right that they cannot have; hence they cannot be liable to the harm of being made liable to be killed. That harm therefore counts in the assessment of whether killing them all is proportionate in the *wide* sense, as it is a harm to which they are not liable. Her claim is then that the aggregate badness of making a thousand people liable to be killed outweighs the wrongful killing of one innocent victim. Killing all thousand is thus disproportionate in the wide sense and cannot be permissible.

In section VI I will deny that killing all would make each liable to be killed. For the moment I will grant her assumption but note that not all harms to which the victims are not liable count in the determination of whether the act that causes them is proportionate in the wide sense. For example, a harm that the victim freely consents to suffer cannot make the act that causes it impermissible by making it widely disproportionate. In the case of the responsible killers, what seems to exclude the costs they incur in being made liable to be killed from the assessment of wide proportionality is that their being made liable consists solely in the withdrawal of a moral shield from liability to which they have no entitlement. The removal of an impediment to the circumstantial conditions of liability is not a relevant cost when the person shielded already satisfies the agential conditions.

Suppose that only one responsible killer is present and that he is about to kill the innocent victim. Another 999 responsible killers are on the way and will arrive in just a few moments. Their imminent arrival overdetermines the killing of the victim, thereby shielding the one responsible killer from liability to be killed by blocking the satisfaction of the circumstantial conditions. But suppose that a third party could prevent the 999 from arriving in a way that would be harmless to them. That action would secure the circumstantial conditions for the single responsible killer’s liability to be killed. This would be bad for him. It would not, of course, make it impermissible to prevent the arrival of the others, for there is only this one harm to weigh against the potential harm to the innocent victim,

18. Ibid.

whereas in the original case there are a thousand such harms. Yet the fact that the removal of the one responsible killer's moral shield against liability would be worse for him does not seem to count *at all* against preventing the others from arriving. It is not a cost that weighs against the threatened harm to the innocent victim in the assessment of wide proportionality. Rather, it seems that the third party has a duty to prevent the one responsible killer from being morally shielded from liability and that the fact that the fulfillment of that duty would be worse for the responsible killer is irrelevant.

There seems, indeed, to be no morally significant difference between the harmless removal of the moral shield in this kind of case and the harmless removal of an innocent shield in the more familiar kind of case. Suppose that a single responsible killer is about to kill an innocent victim and that the only way the victim can save herself is to kill the killer. Her only means of doing this, however, would also kill two innocent bystanders as a side effect. (I think it makes no difference whether they have been placed there by the culpable action of others.) The responsible killer is liable to be killed but is shielded morally by the presence of the two bystanders, the killing of whom would be widely disproportionate. If, however, the victim could harmlessly remove the bystanders, thereby making the killing of the responsible killer widely proportionate, the fact that this would be worse for the responsible killer seems not to count at all against the permissibility of removing them. The only difference is that in this case the removal of the innocent shields secures the permissibility of acting on a preexisting liability justification, whereas in the original overdetermination case the removal of the 999 additional responsible killers secures the liability justification for killing the one by establishing the circumstantial conditions of his liability.

That this difference is not significant is suggested by the fact that the overdetermination case would be relevantly like the bystanders case if we were to assume, as many (such as Rodin) do, that effectiveness, necessity, and narrow proportionality are not internal to liability but are instead external constraints on the permissibility of acting on a liability justification. With that assumption, the responsible killer is in both cases liable to be killed and the harmless removal of the moral shields—again in both cases—simply eliminates an external impediment to the permissibility of acting on the liability justification.

One further objection to Gordon-Solmon's argument that I will mention only in passing is that it seems to imply that it is also impermissible to kill the culpable killers. If, as she says, there is no right not to be made liable to be killed, the culpable killers cannot be liable to the harm of being made liable to be killed. A thousand aggregated instances of this harm should thus outweigh the wrongful killing of the innocent victim, so that killing them would be disproportionate in

the wide sense. Gordon-Solmon is aware of this problem but says only that the harms to the culpable killers in being made liable “either aren’t impersonally bad at all, or aren’t bad enough to clear the threshold of additivity” (a reference to a notion I introduced in earlier work). (8) This, however, does not seem correct, as the relevant harms are ones to which they are not liable and it seems arbitrary to suppose that they do not count just because those who would suffer them are culpable rather than minimally responsible. As I argued in the discussion of Rodin, culpability is relevant to narrow proportionality but not to wide proportionality.

V. Combined Justification

Some philosophers claim that even when there is only a single responsible killer, he is not liable to be killed. If that is right, and if it is also right that a single culpable killer *is* liable to be killed, it is then obvious why it is impermissible to kill a thousand responsible killers and also, perhaps, why it is permissible to kill a thousand culpable killers.

There are at least two reasons why one might deny that responsible killers are morally liable to be killed. One is that, while the explanation I gave of the liability of a responsible killer implicitly appeals to a comparative account of liability, many philosophers assume that the best account of liability is noncomparative, or at least that there is a noncomparative limit to the amount of harm to which a person can be liable. According to a comparative account of liability, when it is unavoidable that someone must be harmed and the only possible options each involve only one of the possible victims suffering a harm of a fixed magnitude, the one who is most responsible for the fact that harm is unavoidable is liable to suffer the harm.¹⁹ According to a noncomparative account, by contrast, there is a limit to the degree of defensive harm to which a person can be liable that is set by the agential conditions of liability—that is, by the magnitude of the harm for which he would be responsible and the degree of his responsibility for the threat of that harm.

One might think that a person who is only minimally responsible for a threatened harm cannot be liable to a defensive harm that is as great as the harm he would otherwise cause. Saba Bazargan, for example, proposes the following formula for determining how much harm a threatener can be liable to suffer in defense of his victim. First, identify the amount of harm the threatener will otherwise cause and multiply it by the threatener’s “percentage responsibility” for the

19. Jeff McMahan, “Who Is Morally Liable to Be Killed in War,” *Analysis* 71 (2011), 552.

threat he poses. That yields the amount of harm he would be liable *for* (not *to*) were he fully culpable. Suppose that a responsible killer is 5 percent responsible for a threat of death, which has a numerical value of -100 . He is, on Bazargan's view, liable to the same amount of harm as a threatener who is fully culpable for a threatened harm of -5 . A fully culpable threatener is liable to defensive harm substantially greater than that which he would otherwise cause—for example, ten times greater. On this assumption, the responsible killer who will otherwise inflict a harm of -100 is liable to be harmed up to -50 in defense of his victim.²⁰

The other reason why one might deny that responsible killers are liable to be killed is that liability is sensitive to matters of agency. When a responsible killer threatens an innocent victim, the choice that the victim or a third party faces is between *intentionally killing* one person as a means and *foreseeably allowing* another person to be killed. If in general the reason not to kill a person is stronger than the reason not to allow a person to die, and if the reason not to harm a person as a means is stronger than the reason not to harm a person as a foreseeable side effect, then there is a moral presumption against killing a person in self-defense. Only if there is a sufficiently significant moral difference between the threatener and the victim can that presumption be overridden.²¹ In most instances, there is such a difference. But one might think that the difference between the responsible killer and the victim—the difference between minimal responsibility and no responsibility—is too slight to overcome the presumption against intentional killing.

These two reasons are not mutually exclusive but may reinforce each other. Suppose, then, that when there is only one responsible killer, he is not liable to be killed. It does not follow that it is impermissible to kill him. It might be permissible to kill him if he were liable to some significant proportion of the harm of being killed and there were a different justification for the infliction of the remainder of that harm. There might, for example, be a lesser evil justification—that is, a justification for the infliction of harm to which the victim is not liable when the only alternative would involve much greater harm to which the victim or victims would

20. Saba Bazargan, "Killing Minimally Responsible Threats," *Ethics* 125 (2014): 114–36. Bazargan writes on p. 121 that "the degree of harm to which an individual is liable is the harm that she threatens to cause multiplied by the percentage degree of her moral responsibility." But the subsequent text indicates that this is not in fact his view. In this sentence he seems to mean "for which" rather than "to which." I am grateful to Matthew Oliver for helpful discussion of Bazargan's text.

21. Jeff McMahan, "Self-Defense and the Problem of the Innocent Attacker," *Ethics* 104 (1994): 252–90, section 1.

also not be liable. I have elsewhere referred to this as a “combined justification.”²² That there can be a combined justification is implicit in the idea that it can be permissible to inflict harm that is disproportionate in the narrow sense provided that the harm to the victim beyond that to which she is liable is proportionate in the wide sense.

Suppose there can be a combined justification for killing a single responsible killer. There is a simple explanation of why that justification is not endlessly repeatable. Imagine a series of cases, in each of which there is one more responsible killer than in the case preceding it. Assume that each responsible killer is liable to the same degree of harm. With each additional responsible killer, there is more harm that has to be justified as the lesser evil. But as the harm that has to be justified increases, the harm to be averted—the death of the victim—remains constant. At some point, the combined harm to the responsible killers beyond that to which they are liable will no longer be the lesser evil in relation to the death of the victim. At that point, the lesser evil element of the combined justification ceases to apply. On the reasonable assumption that there is no other justification for killing further responsible killers, the limit of justification has been reached and the killing of any further responsible killers would be unjustified. This explains why there is a limit to the number of responsible killers it can be permissible to kill in defense of a single victim. If it were known at the outset that the number of responsible killers exceeds the number that could permissibly be killed on the basis of a combined justification, it would be impermissible to kill even one, as that would be ineffective, assuming the agent would respect the limits of the combined justification.

The appeal either to a noncomparative account of the limit of liability or to a strong moral asymmetry between intentional harming and foreseeably allowing harm to occur, together with an appeal to a combined justification, thus seems to offer a good explanation of common intuitions about the killing of responsible killers. Although I have been tempted by this explanation, I now think it is unsatisfactory. While there may be a combined justification for killing one or more responsible killers in certain cases, in most cases there is unlikely to be a combined justification for killing even a single responsible killer.

Suppose that both killer and victim would suffer the same harm in being killed: –100. And suppose that, in accordance with the assumption that a minimally

22. Jeff McMahan, “What Rights May Be Defended by Means of War?” in *The Morality of Defensive War*, ed. Cécile Fabre and Seth Lazar (Oxford: Oxford University Press, 2014), 133–35. For an earlier discussion, see Jeff McMahan, “Targeted Killing: Murder, Combat, or Law Enforcement?” in *Targeted Killings: Law and Morality in an Asymmetrical World*, ed. Andrew Altman, Claire Finkelstein, and Jens David Ohlin (New York: Oxford University Press, 2012), 138.

responsible threatener cannot be liable to a harm as great as that which he would cause, a responsible killer is liable to a harm of no more than -90 (an assumption significantly more favorable to a combined justification than Bazargan's). For it to be permissible to kill the responsible killer, there must be a lesser evil justification for inflicting on the responsible killer the remaining harm of -10 , which is a harm to which he is not liable. That harm must be widely proportionate in relation to allowing the victim to suffer a harm of -100 , to which she is also not liable.

One concern about this reasoning is that it may involve an objectionable form of double counting.²³ The first claim in the combined justification is that the responsible killer is liable to a harm of -90 as a means of preventing the victim from suffering a harm of -100 . The second is that inflicting -10 (which together with the -90 constitutes the harm of death) is the lesser evil in relation to allowing the victim to suffer a harm of -100 . The prevention of the full harm of -100 to the victim is thus serving twice—once to justify part of the total harm to the victim and then again to justify the other part. But the killing of the victim is not being prevented twice. If the prevention of the killing justifies inflicting a harm of only -90 on grounds of liability, that exhausts its power to justify the harming of the killer. If defense of the victim could justify inflicting greater harm on the threatener, he would be liable to greater harm.

There may be a plausible response to this objection.²⁴ Suppose that to defend herself from a responsible killer, a potential victim must harm the killer as a means and harm an innocent bystander as a side effect. She could incapacitate the killer by inflicting a harm of -90 on him but her act would unavoidably inflict a harm of -10 on an innocent bystander as a side effect. Call this the “three-person case.” Suppose that the responsible killer is liable to a harm of -90 as a means of preventing the death of the victim (a harm of -100) and that there is a lesser evil justification for the infliction of a harm of -10 on a bystander as a side effect of preventing the harm to the victim. The defensive act is thus proportionate in both the narrow and wide senses. The harm to the victim is averted only once but it seems to have counted twice – yet this seems legitimate. And it seems that the same can be said of the application of a combined justification to the original case involving only the responsible killer and the victim (the “two-person case”).

In the two-person case, to determine whether killing the responsible killer is proportionate in the narrow sense, one must compare the harm of -90 that the victim would inflict on the responsible killer with the harm of -100 to herself that

23. I am grateful to Stephen Bero and Victor Tadros, who both independently pressed this objection.

24. I am indebted to Patrick Tomlin for this response.

she would avert. To determine whether inflicting harm on the killer beyond that to which he is liable is proportionate in the wide sense, one must compare inflicting the further harm of -10 with the harm of -100 that would be averted. As in the three-person case, the prevention of the full harm of -100 is counted twice, once in the assessment of narrow proportionality and again in the assessment of wide proportionality. The main difference is that in the original two-person case, the harm that must be assessed for narrow proportionality and that which must be assessed for wide proportionality are harms to the same person rather than harms to different people.

I will not attempt to determine whether this comparison is sufficient to rebut the double-counting objection. For it helps us to see what I think is a more serious objection to the application of a combined justification to the original two-person case. There are two significant differences between the two cases.

First, in the three-person case the harm to the responsible killer is completely effective on its own. Together with the assumption that this harm is narrowly proportionate because it is less than the harm it averts, this ensures that the killer is liable to suffer it. In the two-person case, by contrast, the maximum harm to which the responsible killer is potentially liable is -90 . But only death (-100) can be effective; hence the infliction of -90 on its own is wholly ineffective.

Second, in the three-person case, the bystander is harmed as a side effect, whereas in the two-person case, the additional harm of -10 that must be justified as the lesser evil is inflicted as a means. Killing the responsible killer is the means of saving the victim and all portions of the harm of death are included within that means.

Both these facts about the two-person case are problematic for the combined justification. That the infliction of -90 on the responsible killer would be wholly ineffective might suggest that he cannot be liable to it, if effectiveness is internal to liability. One can claim, however, that the single defensive act is overall effective, even if the percentage of the defensive harm to which the responsible killer is liable (ninety percent) would be ineffective if it could somehow be inflicted on its own. Yet if we assume compliance with morality, so that the defensive act will be done only if it is permissible, then the infliction of the -90 can be part of an overall effective act only if the infliction of the other -10 can be justified as the lesser evil. If the infliction of the additional -10 beyond the responsible killer's liability cannot be justified as the lesser evil, the act may not be done and the infliction of -90 in any other way will be ineffective, so that the killer cannot be liable to it.

The harm of -10 cannot, however, be justified as the lesser evil. For there to be a lesser evil justification for the infliction of harm on a person who is not liable to that harm, the harm caused must be sufficiently small, or that prevented sufficiently

great, to justify overriding the constraint against harming. A harm of -10 is ten percent of the harm of death. Suppose that in both the two-person and three-person cases, both the victim and the responsible killer would lose fifty years of good life in being killed. On that assumption, ten percent of the harm of death is equivalent to the loss of five years of good life. In the three-person case, it would not be justifiable as the lesser evil to cause an innocent bystander a loss equivalent to the loss of five years of good life as a side effect of preventing the innocent victim from losing fifty years of good life. That would be disproportionate in the wide sense.

In the two-person case, effective defense requires that the same harm be inflicted not as a side effect but as an intended means. Most discussions of wide proportionality presuppose that it is sensitive to the distinction between means and side effect. They presuppose, that is, that to offset an intended harm to which the victim is not liable, it is necessary to prevent a harm greater than that which would be necessary to offset an equivalent unintended harm.²⁵ If this familiar assumption is correct, then the harm beyond that to which he is liable that it might be widely proportionate to inflict on the responsible killer must be *less* than that which it would be widely proportionate to inflict on the innocent bystander in the three-person case.

In the three-person case, it seems plausible to suppose that the harm it might be widely proportionate to inflict on an innocent bystander as a side effect of saving the innocent victim must be less than -1 —that is, less than one percent of the harm of death. If death involves the loss of fifty years of good life, one percent of that is a harm equivalent to the loss of a half a year of good life. If harms inflicted as a means are harder to justify than equivalent harms inflicted as a side effect, the harm it might be widely proportionate to inflict on the responsible killer in the two-person case beyond that to which he is liable must be significantly less than -1 .

One might argue that because the harm that must be justified as the lesser evil in the two-person case would be inflicted on a person who is already liable to a substantial harm, it can be greater, other things being equal, than the maximum that could be justified if the victim were not liable to any harm at all, as is true of the innocent bystander in the three-person case. This would be analogous to the view that to inflict a certain amount of punishment on a guilty person beyond what he deserves is less objectionable than to inflict the same amount of punishment on an innocent person. I am skeptical of both these views but cannot discuss them here.

25. "Proportionate Defense," section VII.

Putting aside this last point, it seems reasonable to conclude that in the two-person case, the amount of harm that it can be justifiable to inflict on the responsible killer *as the lesser evil* is substantially less than -1 . That means that for there to be a combined justification for killing the responsible killer, he must be liable to a harm greater than -99 —that is, greater than ninety-nine percent of the harm of death. While this is formally compatible with the assumption that a responsible killer cannot be liable to as much harm as he would cause, in substance it trivializes that assumption. It seems, therefore, that there cannot be a combined justification for killing even a single responsible killer. Yet very few people accept that it is impermissible for an innocent victim to kill a responsible killer in self-defense. Those who deny that a responsible killer is liable to be killed have therefore been assuming that some form of combined justification can provide the required justification. But the foregoing discussion shows, I think, that this is mistake.

VI. Proportionality in the Aggregate

The problems I have been discussing—how, if there is a liability justification for killing one, there can also be a limit to the number it can be permissible to kill, and why the number of killers seems more obviously constraining when they are culpable than when they are minimally responsible—have thus far proved intractable. In earlier work, I proposed a solution about which I now have doubts.²⁶ I will not review that suggestion here but will instead pursue a new and different line of thought that seems more promising.

Assume that if there is only one responsible killer, he is morally liable to be killed, provided that the circumstantial conditions are satisfied. Now suppose that there are 999 others. The killing of the innocent victim is now overdetermined, which means that killing the original responsible killer, or any one of the others, has become either ineffective (if life during the interval before the next killer would act would not be worth living) or narrowly disproportionate (if that interval would be worth living, though short). If effectiveness and narrow proportionality are circumstantial conditions of liability rather than external conditions of permissibility, none of the thousand responsible killers is liable to be killed. Each has a moral shield against liability provided by the presence of the others. But, as I noted earlier, none has a right to this shield. If there were only one responsible killer and a third party could prevent the other 999 from being present without harming them, she would have a duty to do that. Similarly, if there were a thousand responsible killers and the third party could remove 999

26. McMahan, "Who Is Morally Liable," 554–55; and McMahan, "Duty, Obedience, Desert," 155–57.

without harming them, she would have a duty to do that. But suppose that the only way that anyone can remove them is to kill them. Only if all are killed can the victim survive. Yet if all are killed, each killing is effective in preventing the victim from losing a significant amount of good life. For in that case each killing occurs in a context in which no one of the other 999 responsible killers will kill the victim. The killing of each responsible killer is morally as it would be if he were the only one present. In short, killing all the responsible killers establishes the circumstantial conditions for the liability of each. Yet, assuming compliance with morality, all will be killed only if it is permissible to kill them.

This again raises the question whether the permissibility or impermissibility of killing all can be determined independently of determining whether each is individually liable to be killed. I suggest that, when we consider the harm that would be involved in killing a thousand responsible killers, taking into account that each is only minimally responsible for the threat he poses, and compare that harm to the threatened harm to the victim, taking into account that she bears no responsibility for the fact that killing is unavoidable in the circumstances, we can judge intuitively that killing all thousand would be disproportionate. We can see this without first determining whether each of the thousand responsible killers is liable to be killed. This is a judgment of proportionality that, like narrow proportionality but unlike wide proportionality, takes account of the moral responsibility of those who would be harmed. But, like wide proportionality and unlike narrow proportionality, it also takes account of the number who would be harmed. We can refer to this form of proportionality as *proportionality in the aggregate*.

I believe that we can infer that killing all thousand responsible killers is disproportionate in the aggregate just from facts about the harms that would be caused and prevented, the varying degrees of responsibility of the different possible victims, and the number of people who would be killed in each option. If this is right, it resolves the circularity problem mentioned in section II. That was that it seemed that whether it is permissible to kill all thousand responsible killers depends on whether each is liable, which depends on whether killing each would be effective, which in turn depends on whether the others will be killed, which depends, assuming compliance with morality, on whether it is permissible to kill them all. But whether it is permissible to kill all thousand does *not* depend on whether each is liable. We can know that it is *impermissible*, because disproportionate in the aggregate, without knowing whether each is liable.

When they consider the culpable killers, many people—perhaps most—think it is permissible to kill all thousand. If so, they implicitly judge that killing all is proportionate in the aggregate. If that is so, and if the victim has the ability to kill all the culpable killers sequentially, it is then permissible for her to kill each in

turn. Because in killing each she knows that it will be permissible for her to kill the others, she is justified in regarding each killing as effective and narrowly proportionate. Each killing is narrowly proportionate because it affords her the opportunity to kill the remaining culpable killers in the sequence, thereby enabling herself to have many further years of good life.

My claims about the responsible killers may seem to have a rather paradoxical implication. I claim that it is impermissible to kill all the responsible killers because killing them all would be disproportionate in the aggregate. Yet if the victim does kill them all, she thereby creates the circumstantial conditions of the liability of each. It seems, in other words, that in killing them all, she establishes the conditions of a liability justification for killing each, thereby establishing the permissibility of killing each, thus establishing the permissibility of killing all. (This is the reasoning that Gordon-Solmon seeks to block with her claim that making all thousand liable is widely disproportionate.)

Suppose that the thousand responsible killers are queued up at a distance and will be appearing one by one, in rapid succession. In this version of the example, however, the victim does not have the means to kill them one at a time. But she has a large artillery weapon that can kill all of them at once. If she fires this weapon, it may seem that her act creates the conditions of its own permissibility; for by simultaneously killing them all, it ensures the satisfaction of the circumstantial conditions of liability of each person killed, thereby establishing a liability justification for each killing.

I think, however, that we should not accept that the act guarantees its own permissibility. The act of killing all is impermissible because it is disproportionate in the aggregate. If the victim fails to comply with morality and kills all thousand, she will have done what was disproportionate in the aggregate and her act will still be objectionable on that ground. It is true that she will also have created the circumstantial conditions of the liability of each victim. But those conditions did not obtain when she acted. The thousand responsible killers were not liable to be killed when she acted, and it was impermissible for her to create the conditions of their liability by killing them. She could not, therefore, act in the expectation of permissibly creating the conditions of individual liability, as she *could* if the killers were culpable.

We should accept, however, that if she does kill them all, she makes it the case that none of them has been *wronged* by being killed. This should not be surprising, for in killing all she has ensured that each killing was both effective and narrowly proportionate. Suppose the victim had had a gun and 999 bullets and had killed the first 999 responsible killers, only to be killed by the last one. In that case, each of the 999 killed *would* have been wronged, for he would have been killed wholly gratuitously. But if she kills all thousand, none of them can complain

of having been killed gratuitously, for killing him was, in the circumstances, necessary and effective in securing her survival.

But, again, that in killing them all she makes it the case that none was wronged does not entail that she acted permissibly, on the basis of a set of liability justifications. The reason her act was impermissible—that it was disproportionate in the aggregate—is a matter of the numbers, which is independent of whether any individual is wronged.

Earlier I conceded that the view I would defend is similar to Rodin's view. I will close by explaining how my appeal to the notion of proportionality in the aggregate differs from his view that there is a lesser evil obligation not to act on the liability justifications for killing the responsible killers. The main difference is that he believes that all the responsible killers are liable to be killed because they all satisfy the agential conditions of liability, which are, in his view, the only conditions of liability. Even when there are a thousand responsible killers, there is a liability justification for killing each. (This implies—implausibly, in my view—that if the victim kills only one of the thousand, she does not wrong him, for he has forfeited his right not to be killed, even gratuitously.) If it is impermissible to act on those justifications, it seems that they must be overridden by countervailing considerations. He claims that what overrides the justifications are the harms that the killers would suffer in being killed. I claimed that this implies that the justifications are overridden by the same harms that they justify.

My view, by contrast, denies that there are liability justifications for killing any of the thousand responsible killers. Even though killing them all would secure the circumstantial conditions of their liability, killing them all is impermissible because it would be disproportionate in the aggregate. And even if the victim or a third party were to kill them all, that would not justify their having been killed by retroactively supplying a liability justification for each killing. It would merely make it the case that none of those who had been wrongly killed had been wronged in being killed. The wrongness of killing them was not in any wrong done to each but was instead a matter of the numbers.²⁷

27. I am very grateful to the Institute of Advanced Study at the University of Birmingham for generous support during the writing of this essay. I have greatly benefited from comments on earlier versions by Ben Bronner, Tony Coady, Roger Crisp, Adam Gastineau, Kera Gordon-Solmon, Richard Holton, Andrew Lister, Victor Tadros, Larry Temkin, Jesse Tomalty, Patrick Tomlin, Stephen Woodside, and, especially, Derek Parfit.