

10

The Limits of Self-Defense

JEFF MCMAHAN

1 Necessity Does Not Require the Infliction of the Least Harm

According to the traditional understanding of necessity in self-defense, a defensive act is unnecessary, and therefore wrong, only if there is an alternative means of defense with at least an equal probability of success that would cause less harm.¹ There are obvious counterexamples to this understanding of necessity. For example:

Counterexample 1

An innocent victim who is threatened with death by a culpable threatener has only two defensive options. Both are certain to be effective.

- (1) She could incapacitate the threatener by breaking his arm, or
- (2) she could incapacitate him in a way that would not harm him but would break the finger of an innocent bystander as a side effect.

Although the second option would cause less harm, the victim ought to choose the first option instead. This is because the threatener has made himself morally liable to have his arm broken while the innocent bystander would be

¹ The material in this subsection draws on lectures I presented at various universities over a number of years. One version of the text on which these lectures were based is available at http://iis-db.stanford.edu/evnts/6657/McMahan_ProportionalityRevised_5-2011.pdf. The relevant pages are 21–24. There is also some overlap between what I say in this section and Seth Lazar, “Necessity in Self-Defense and War,” *Philosophy and Public Affairs* 40 (2012): 3–44, though his very fine and wholly original article goes into greater detail and covers many issues I do not attempt to discuss. In our thinking about necessity, Lazar and I have worked independently of one another.

wronged by having her finger broken. Following Seth Lazar, we might say that the first option would cause less “morally weighted” harm.²

More importantly, there are cases in which the better defensive option has a lower probability of success. For example:

Counterexample 2

An innocent victim who is threatened with death by two culpable threateners has only two defensive options. (1) She could kill the two culpable threateners. This would give her a 99 percent probability of saving her life but, even if successful, would not prevent the threateners from inflicting a minor harm on her. (2) She could act in a way that would incapacitate the culpable threateners without harming them. This would give her a 100 percent probability of avoiding being harmed in any way but would kill an innocent bystander as a side effect.

In this example, the victim’s second option has a higher probability of successful defense, offers a more complete defense, would cause less harm, and the harm caused would be a side effect rather than, as in the first option, an intended means. According to the traditional interpretation of necessity, option two satisfies the necessity condition because there is no alternative means of defense that has an equal or higher probability of averting the threatened harm. Option 1, which has a lower probability of success, therefore seems ruled out as unnecessary. Yet it is clear that the potential victim must choose option 1. She may not increase her probability of survival by 1 percent and avoid a minor harm at the cost of killing an innocent bystander.

It may seem that the obvious explanation of this is that the two culpable threateners are morally liable to be killed, whereas the innocent bystander is not. But this is only one element of the deeper explanation, which is that option 1 is proportionate while option 2 is not. Both culpable threateners are liable to be killed in defense of their potential victim and thus killing both of them would be proportionate in the sense (which I call the “narrow” sense) that the harm would not exceed that to which they are liable. The innocent bystander is of course not liable to any harm, so for it to be proportionate to kill him as a side effect of self-defensive action, the harm he would suffer in dying must be justifiable as the “lesser evil”—that is, it must be substantially less than the harm the act would prevent. But it is reasonable to suppose that that is untrue—that is, that death would not be substantially less bad for

² Lazar, “Necessity in Self-Defense and War,” 7.

the bystander than for the potential victim. On that assumption, killing the bystander as a side effect would be disproportionate in what I call the “wide” sense (in which proportionality is a constraint on a lesser-evil justification). It is because a proportionate option is always morally better than a disproportionate one that option 1 is better than option 2.

There are, however, counterexamples to the traditional understanding of necessity in which all the options are proportionate. For example:

Counterexample 3

There are two means of preventing 100 innocent people from being killed by ten culpable threateners. (1) If a third party were to kill all ten culpable threateners, that would have a 99 percent probability of saving all 100 victims. (2) If a third party were to kill one culpable threatener as a means, that would have a 100 percent probability of saving all 100 victims but would also kill two innocent bystanders as a side effect.

By threatening to kill 100 innocent victims, the ten culpable threateners seem to make themselves liable to be killed as a means of saving the victims (albeit a means with only a 99 percent probability of success); therefore killing all ten of them would be proportionate in the narrow sense, as, a fortiori, killing only one of them would be. It is reasonable, moreover, to suppose that killing two innocent bystanders is proportionate as a side effect of preventing (via the killing of one person who is liable to be killed) 100 different innocent people from being wrongly killed. So both options seem proportionate. As in Counterexample 2, option 2 has a higher probability of successful defense and would cause less harm (in that it would involve killing three people rather than ten). It therefore seems to satisfy both the proportionality condition and the necessity condition as traditionally interpreted. Because option 1 has a lower probability of success, it seems to be ruled out by the traditional interpretation of the necessity condition.

This, however, is a perverse implication of the condition that evaluates a defensive option by comparing it with alternative means of avoiding a threatened harm. Such a condition ought not to exclude, but ought instead to require, the defensive option that is *morally best*, all things considered. And in determining which defensive option is morally best, this condition must allow for trade-offs between the probability that the defensive action will succeed (as well as the degree to which the defense will be complete) and the expected harm that the defensive action will cause. That is, the morally best and therefore necessary defensive option is the one that, in comparison with other options that have different probabilities of success and would cause different

amounts of harm, has the morally best trade-off between the harm prevented and the harm caused. Call this the *trade-off interpretation* of the requirement of necessity. It contrasts with the traditional interpretation, which requires the choice of the option that would cause the least harm only among the defensive options that have the highest probability of successful defense. According to the traditional interpretation, while the defensive option that would be most effective in terms of probability and completeness may be ruled out by proportionality, it cannot be ruled out on grounds of necessity.

But, again, it is a mistake to suppose that the condition that evaluates a means of defense by comparison with alternative means of avoiding the threatened harm can insist that the morally best means must have the highest probability of success. In Counterexample 3, for example, the third party seems morally required to accept a slightly lower probability of success in attempting to defend the 100 innocent victims rather than ~~killing~~ two innocent bystanders, even when killing the two innocent bystanders would be proportionate in relation to the saving of 100 innocent people. Assuming this is correct, one must conclude that option 1 satisfies the necessity condition while option 2 does not.³

An explanation is still required of why option 2 is morally worse, and therefore ruled out as unnecessary, even though it offers the best chance of successful defense. One might think that the relevant difference between the two proportionate options in Counterexample 3 is that the killings in option 1 would have a liability justification whereas two of the three killings in option 2 would have only a lesser-evil justification. If one also thought that a liability justification is a better or stronger form of justification than a lesser-evil justification, and thus has priority when a person can act on a justification of either sort, one might think that this explains why the victim in Counterexample 3 ought to choose option 1 rather than option 2.

Yet it seems that a liability justification may not always have priority over a lesser-evil justification. Consider, for example:

Kill or Bruise

Innocent Victim will be killed by highly culpable Threatener unless she either (1) kills Threatener, or (2) acts in a way that will incapacitate him without harming him but will cause innocent Bystander to suffer a painful bruise as a side effect.

³ In the lecture text cited in note 1, I suggest that what I here call the trade-off interpretation be understood as a further principle in addition to the traditional requirement of necessity. I now think that was a mistake.

Intuitively it seems clear that Victim ought to bruise Bystander rather than kill Threatener, yet it also seems that there is a liability justification for killing Threatener but only a lesser-evil justification for bruising Bystander.

It may be, however, that in *Kill or Bruise*, there is in fact no liability justification for killing Threatener. This is true if three claims are correct. These are (1) that the trade-off interpretation of the necessity condition is the right interpretation, (2) that bruising Bystander is morally the better of the two defensive options, so that killing Threatener is unnecessary according to the trade-off interpretation, and (3) that necessity is “internal” to liability, in the sense that a person cannot be liable to an unnecessary harm. Whether necessity is internal to liability is, however, the topic of ~~the next subsection~~, so I will defer consideration of this issue till then.

2 Converting Assessments of Necessity to Assessments of Proportionality

We still have not determined why killing the culpable threateners, or ~~more~~ more culpable threateners, is morally better in Counterexamples 1, 2, and 3 but not in *Kill or Bruise*, even though in all four cases the option that involves killing more culpable threateners has an equal or greater probability of success and inflicts less harm overall. There is, I think, no simple answer to this question, though there is a simple method one can follow in trying to determine which option is better in particular cases. The method is perhaps best illustrated by a variant of *Kill or Bruise*.

Kill or Allow Bruising

Innocent Victim will be killed by highly culpable Threatener unless she either (1) kills Threatener, in which case she will be completely unharmed, or (2) incapacitates him without harming him, in which case his attempt to kill her will cause her a painful bruise.

A natural way to think about this example is to ignore the option of submission, which is clearly not required, and to treat the two defensive options as exhaustive of the possibilities. One can then treat the option involving the lesser harm as the default option (that is, what *will* be done unless the first option is adopted instead) and ask whether the infliction of the greater harm would be *proportionate* as a means or side effect of avoiding the lesser harm. This has the form of a proportionality judgment when the option of not inflicting the

greater harm is understood as tantamount to submission to the lesser harm. Thus, because incapacitating Threatener at the cost of being bruised is clearly permissible, the only justification there could be for killing Threatener is to prevent him from bruising Victim. It would, however, be disproportionate to kill even a highly culpable threatener (in this case, one who is *attempting* to commit murder) as a means of preventing him from merely causing his victim a painful bruise. In *Kill or Allow Bruising*, therefore, the option of incapacitating Threatener is morally better, so that killing him is ruled out as unnecessary. (One could apply the same method to *Kill or Bruise*. In that case killing Threatener would be disproportionate as an alternative to the default option of bruising Bystander.)

Suppose that, rather than incapacitating Threatener in a way that would allow him to bruise her, Victim's second option was to prevent him from killing her in such a way that he would still be able to break both her legs, which she knows he would do. Again, we could convert this to an issue of proportionality by ignoring the option of submission and treating the breaking of her legs as the default. We could then ask whether, when a person is culpably attempting to kill her, it would be proportionate for her to kill him as a means of preventing him from breaking both her legs. If it would be, then in the original case with three options we could conclude that Victim has no option that is morally better than killing Threatener, so that killing him would be necessary.

It is perhaps worth noting that this case shows that an option's being necessary in the relevant sense does not entail that no other option is permissible, at least for certain agents. While it would be permissible for Victim to kill Threatener, it would also be permissible for her to incapacitate him in a way that would allow him to break her legs. I suspect, however, that it would be impermissible for a third party to incapacitate rather than kill Threatener, thereby allowing him to break Victim's legs, if the third party had a duty to defend Victim and neither option were more costly to her than the other. But I leave this question open.

This way of thinking about necessity must of course be extendable to cases involving more than two options for the *evasion* of a threatened harm. The way to do this might be to identify the option that would involve the least harm, or perhaps the least harm to which the victims would not be liable, and treat that as the default. One might then compare that option with each of the others, one at a time. Any option that is disproportionate as an alternative to the default option involving the least harm would be ruled out as unnecessary. Any option that is morally better than the least harmful one would replace it as the default option in the remaining pairwise comparisons. In this way

one would eventually reach the best option, or equal best options. All other options are unnecessary and therefore impermissible.

There is, however, a powerful objection to supposing that judgments of necessity can be made in this way. It may be that option A is better than option B and option B better than option C, without it following that A is better than C; for there may be factors relevant to the comparison between A and C that are not relevant in the previous pairwise comparisons. Because of this, it is possible that C could be better than A.⁴ If, therefore, we rule out C because it is morally worse than B and then conclude that only A is necessary, we will have made a mistake. This problem seems, however, quite general and may present no more of a challenge to what I have suggested than it does to any other method of determining which among a range of defensive options is morally best. But if there is a method that escapes this challenge, it is likely to be preferable.

3 Effectiveness Is Internal to Liability

As I mentioned in the previous section, one important question about necessity in defensive action is whether it is internal to liability itself, in the sense that a person can be liable to be harmed in a certain way only if harming him in that way is necessary in the relevant sense—that is, according to the trade-off interpretation, only if harming him in that way is the morally best means of achieving a legitimate defensive aim, or an unavoidable side effect of that means of achieving the aim. I noted, for example, that in *Kill or Bruise*, if the option of killing Threatener fails to satisfy the trade-off interpretation of the necessity condition because it is morally worse, all things considered, than bruising Bystander, and if necessity is internal to liability, then Threatener is not liable to be killed. This implication is rejected by many philosophers who write on the ethics of self-defense. While some of these philosophers seem to accept the traditional interpretation of the necessity condition, the reason why most reject the implication that Threatener is not liable is that they explicitly reject the claim that necessity is internal to liability.

I have argued in the past that necessity is internal to liability (henceforth, “necessity is internal”). I will continue to defend that claim here, though I will later acknowledge that there may be exceptions to it. The possibility that there

⁴ The locus classicus for this problem is Larry Temkin’s important book *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (New York: Oxford University Press, 2012).

are exceptions arises only if the trade-off interpretation of the necessity condition is correct. If the traditional interpretation is correct, then my arguments suggest that necessity is always internal to liability. Alternatively, one might say that in cases in which a harm to a threatener is unnecessary in the traditional sense, the threatener cannot be liable to suffer it; yet in some cases in which a harm to a threatener is unnecessary only in the trade-off sense, it is possible that the threatener is nonetheless liable to suffer it.

It is important first to distinguish necessity from effectiveness.⁵ Although a harm that would be wholly ineffective cannot of course be necessary, a defensive harm can be effective yet unnecessary—either because there is an alternative that would either be equally effective but cause less morally weighted harm, or because there is an alternative that would be better in terms of the trade-off between the probability of successful defense and the degree of harm caused.

There are two dimensions of effectiveness, both of which are matters of degree. One is the probability that harm will be averted and the other is the completeness with which harm will be averted. A defensive option might, for example, have a 99 percent probability of completely averting a harm, or it might have a 100 percent probability of averting 99 percent of the threatened harm. I will in general use “effectiveness” to refer to both probability and completeness.

Because effectiveness is different from necessity, it is a separate question whether effectiveness is internal to liability—that is, whether a person can be liable to a harm that would be ineffective. I believe that it is of considerable importance theoretically that the concept of liability be understood as having an internal effectiveness condition. If effectiveness is internal to liability, liability is then essentially *instrumental*, in the sense that a person cannot be liable to be harmed unless harming him is either a means or an unavoidable side effect of preventing or rectifying a different harm. This then differentiates liability from desert, which is not instrumental. As desert is commonly understood, a person can deserve to be harmed, and there can be a reason to harm him, even if harming him will not prevent or rectify any other harm. There can be a reason to harm him even though all further harm is avoidable. If, by contrast, liability is instrumental, then a person can be liable to be harmed only when some harm is unavoidable. Liability is essentially a matter of justice in the distribution of harm when some harm is unavoidable. The person or persons who are liable are those whom it is most just, or least unjust, to harm in the circumstances. If the harm does not go to them, it will inevitably go to others. This is why someone’s being liable to be harmed constitutes a pro tanto justification for harming him.

⁵ A point emphasized by Helen Frowe, *Defensive Killing* (Oxford: Oxford University Press, 2014), chap. 4.

It is important to be clear about this difference between desert and liability because the bases of desert and the bases of liability may be much the same: namely, voluntary action that risks or causes unjustified harm together with certain mental states that are necessary for moral responsibility or culpability. But liability is morally less contentious because it arises only when someone must be harmed. It is possible, therefore, that harm inflicted on the basis of liability is always regrettable and bad in itself, whereas deserved harm is generally regarded as good in itself, even though it is bad for the person who deserves it.

We can perhaps test for the plausibility of the idea that effectiveness is internal to liability by considering an example—a variant of the familiar trolley case.

Shock

Murderer intends to kill Victim, who is trapped on a branch trolley track. To kill her, Murderer must divert an oncoming trolley from the main track to the branch track. To do this, he must depress a large lever continuously for a number of minutes. Victim has the ability to send one, but only one, powerful electric current along the track that will travel to the lever, delivering a brief but intensely painful shock to Murderer. Murderer knows this and to ensure that he does not release the lever, has strapped himself to it, using the weight of his body to hold it down. If Victim sends the current, the shock will be brief, so that even if Murderer seeks to unstrap himself, he will be unable to do so until after the shock has ceased.

It is certain that Victim's sending the shock will fail to save her. There might nevertheless be three distinct justifying reasons for her to send it: that Murderer ~~might~~ deserve to be harmed, that harming him would constitute an assertion of her moral dignity, and that the harm might have a deterrent effect, either on Murderer himself or on others. But I will bracket these considerations in this discussion. My question is not whether Murderer deserves to be harmed but whether he is liable to it. Let us therefore assume, for the sake of argument, that no one can deserve to be harmed. Still, Murderer may well be liable to be shocked as a means of affirming the Victim's dignity or moral status, and in that case the harm would not be ineffective. Let us therefore also assume, again for the sake of argument, the Gandhian view that sending the shock when it would be ineffective defensively would not in fact affirm but would instead demean Victim's moral dignity. Finally, assume that even without the shock, Murderer will never seriously harm anyone again and that no one else would know about the shock if it were administered, so that it would have no deterrent effects.

Even if we agree to accept these assumptions, it is unlikely that we will be able to prevent the contrary assumptions from affecting our intuitions about the case. ~~That may be sufficient to explain the lingering intuition that it would be permissible for Victim to deliver the shock, and that Murderer would not be wronged by being harmed in this way. If so, it is unnecessary to suppose that Murderer has forfeited his right not to be harmed in a way that would be pointless.~~

There are several reasons why effectiveness must be internal to liability to defensive harm. One is that ineffective harming is wholly gratuitous—unless it is independently justified on other grounds, such as desert, in which case the question remains whether it is also justified on grounds of liability. If liability is a matter of justice in the distribution of harm when harm is unavoidable, then a person cannot be liable to be gratuitously harmed. The standing presumption is that to harm a person is to wrong him. There are numerous forms of justification that can override this presumption. But to harm a person deliberately when he does not deserve it and doing so serves no purpose whatsoever—in short, to harm him without any positive moral justification—is to wrong that person. It is, in other words, to cause him harm to which he is not liable.

A second reason why effectiveness must be internal to liability is that if it were not, it would be impossible to determine whether acts that inflict foreseeably ineffective harms to which people are allegedly liable are either proportionate or necessary. Suppose, for example, that Victim gives Murderer the electric shock. For this harm to be proportionate, it must not be excessive in relation to its good effects. But by hypothesis it has no good effects; it is wholly gratuitous. That may not mean that the act is necessarily disproportionate; it may instead be that the question of proportionality simply does not arise. Much the same is true of necessity in this case. The question whether the electric shock is necessary does not arise unless there is some end to which it might be instrumental.

A defender of the view that effectiveness is not a condition of liability to defensive harming might respond by claiming that whether the shock is proportionate is a matter of how it compares, not to any harm prevented, but to the wrong to which it is a response. The wrong is Murderer's killing of Victim, and clearly an intense shock is proportionate in relation to murder. But it is now unclear how liability is supposed to differ from desert. Desert, as I earlier observed, is not instrumental. It is entirely retrospective. What a person deserves is determined wholly by the nature of what he has done. Hence, to determine whether harm inflicted on grounds of desert is proportionate, one must compare that harm with the wrong committed. But this is not the way to determine whether harm inflicted on grounds of liability to defensive action is proportionate. That is determined by the comparison between the harm inflicted and the wrongful harm to be prevented. (It is worth mentioning here

the notion of liability to harm as *reprisal*. The relevant criteria of proportionality in harm in reprisal are both retrospective and prospective. For a harm to be proportionate as a matter of reprisal, it cannot be significantly greater than the harm to which it is a response and it must not be excessive in relation to its expected deterrent benefits. The same considerations that favor the view that effectiveness is a condition of liability to defensive harm also favor the claim that effectiveness in the achievement of deterrence is a condition of liability to harming in reprisal.)

Liability, in short, is a form of justification that is distinct from desert. What makes it distinctive, both conceptually and morally, is that it is instrumental. For a person to be liable to harm on grounds of defense, harming him must have a defensive effect. Similarly, for a person to be liable to harm as a matter of corrective justice, harming him must be effective as a means of compensation. (And for a person to be liable to harm as a matter of reprisal, as distinct from retribution, harming him must have a deterrent effect.) Particularly for those who reject the idea that people can deserve to be harmed but nevertheless believe that it can sometimes be permissible to harm people precisely because of what they have voluntarily done, liability as an instrumental notion is essential.

4 Necessity Is Internal to Liability

Even if effectiveness is internal to liability to defensive harming, it does not follow that necessity is; for, as I noted earlier, defensive harming can be effective but unnecessary. Here are two examples, one involving the infliction of some harm rather than none, the other involving the infliction of more harm rather than less.

Standing One's Ground

Threatener will culpably kill Victim unless Victim takes some preventive action. She has two options, both of which she knows to have a 100 percent probability of successfully protecting her. She can (1) retreat from the confrontation without risk or cost and alert the police, who will then be able to subdue and constrain Threatener without harming him, or (2) stand her ground and kill him in self-defense.⁶

⁶ Helen Frowe discusses a parallel though more fanciful example in Frowe, *Defensive Killing*, chap. 4.

Second Leg

Threatener will culpably kill Victim unless Victim engages in self-defense. She has two options. She has a gun and is highly skilled in using it. But she has only one bullet. As Threatener rushes toward her with a meat cleaver, she can either (1) shoot him in one leg, or (2) shoot him when one leg is behind the other, so that her bullet will pass through one leg and hit the other. Either option would be certain to incapacitate him.

In *Standing One's Ground*, killing Threatener would be wholly effective but is unnecessary on any plausible interpretation of the necessity constraint, as Victim can protect herself equally effectively in a way that would cause no harm to Threatener (other than what he would suffer in later being prevented from harming others). Many philosophers nevertheless believe that by wrongly threatening Victim's life, Threatener forfeits his right not to be killed, and hence would not be wronged if Victim were to kill him. On their view, while effectiveness is internal to liability, necessity is not.

In *Second Leg*, wounding Threatener's second leg is unnecessary. This case is interesting because it may be unclear whether Victim's second option counts as effective. While the *act* of shooting Threatener in both legs is effective, the *harm* caused to the second leg is ineffective, since it achieves nothing beyond what is already achieved by the wound to the first leg. I will not, however, pursue this complication here but will treat *Second Leg* simply as a case of two effective options, one of which is unnecessary. Those philosophers who claim that necessity is not internal to liability believe that Threatener is not wronged if Victim shoots him in both legs rather than in only one. They accept that it is wrong to shoot him in both legs, but claim that that is not because the unnecessary harm wrongs him or violates his rights.

It seems to me, however, that Threatener is clearly wronged by being killed in *Standing One's Ground* because the unnecessary harm is entirely gratuitous and no one can be liable to suffer gratuitous harm. The same is true of shooting Threatener through both legs in *Second Leg*. There is, moreover, a further reason for thinking that Threatener is wronged by being shot through both legs in *Second Leg*. In this case, Victim waits for the right moment so that her one bullet will strike both of Threatener's legs rather than just one. Imagine a variant of the case in which Victim is unable to strike both legs with one shot but has two bullets rather than one. She shoots Threatener in one leg with the first bullet, thereby incapacitating him, and then immediately shoots him in the other leg with the second bullet. It is clear that in firing the second shot she wrongs Threatener. But there seems to be no morally significant difference between her action in this case and her action in *Second Leg*. This is particularly

problematic for those who accept that effectiveness is internal to liability but deny that necessity is. For in the variant case the second shot is necessarily ineffective, as the defensive aim has already been achieved. (This parallels the fact, which I noted earlier, that the *harm* to the second leg in the original case is ineffective, even if the act is effective.) So the defender of this pair of claims seems to have to accept that Threatener is wronged by the second shot in the variant case but is not wronged by being shot through both legs in the original case. But that is implausible.

I have argued that the view that necessity is external to liability has implausible implications. But those who hold that view have argued that the view that necessity is *internal* has implausible implications. A defense of this latter view must address these objections.

4.1 OBJECTION 1: IF NECESSITY IS INTERNAL, THREATENER WOULD BE WRONGED

One implication that many people find implausible has already been mentioned—namely, that a potential murderer, such as Threatener in *Standing One's Ground*, would be wronged by being killed in effective though unnecessary self-defense by his potential victim. I have already responded to this objection by arguing that no one can be liable to be killed gratuitously. I concede that in most actual cases there are reasons why killing that is effective but unnecessary for physical defense is not gratuitous. For example, the threatener may deserve to be harmed and may be liable to suffer a certain amount of harm as a means both of preserving deterrence and of affirming the victim's moral status. But I want to put these considerations aside and focus on the question whether a threatener is wronged by the infliction of defensive harm that is effective but unnecessary for the achievement of successful defense.

4.2 OBJECTION 2: IF NECESSITY IS INTERNAL, THREATENERS ARE WRONGED TO THE SAME DEGREE

A second, related objection to the claim that necessity is internal is that, even if threateners are wronged by being harmed unnecessarily, they are not all wronged to the same degree. Suppose, for example, that in *Standing One's Ground*, Threatener is attempting to kill Victim only because he has unwillingly been given a hallucinogenic drug and believes that Victim is trying to kill him. Victim knows this but kills Threatener rather than fleeing to safety. The view that necessity is internal implies that Threatener is wronged by being killed in this version of the case, as well as in the original version. But it seems

clear that he is wronged *to a greater degree* in this second version. And it has been claimed that the view that necessity is internal “will find it very hard to explain this difference, since in both cases the person who is [killed] is not liable to be killed . . . as the killings do not meet the necessity condition, and therefore both acts must constitute equal wrongs on this view.”⁷

The response to this is that it is compatible with the view that necessity is internal that different people can be wronged to different degrees by being caused to suffer equivalent harms to which they are not liable. Many think that this can be true because of facts about the agent. They think, for example, that one victim is wronged to a greater degree than another if the agent harms the one as an intended means but harms the other to the same degree as an unintended side effect. But it can also be true because of facts about the victims. I will cite three types of example.

First, suppose that Drivers 1 and 2 are driving recklessly late at night. Both their cars repeatedly veer off the road, but the terrain is flat and they are able to regain the road without crashing. Driver 1 runs over and kills Victim 1, who had decided to sleep on the shoulder of the road. Driver 2 hits and kills Victim 2, who is walking well away from the edge of the road. Because Victim 1 had assumed a risk that Victim 2 had not, he is wronged to a lesser degree than Victim 2, even though neither is liable to be harmed in any way.

Second, suppose that it is extremely morally important to act in a way that will unavoidably inflict a certain amount of harm, $x + n$, on P_1 and P_2 . P_2 is not liable to suffer any harm but, because of her wrongdoing, P_1 is liable to suffer harm of amount x , which has been justifiably inflicted on her. The remaining harm, n , can be inflicted on either P_1 or P_2 . Neither is liable to this harm and so would be wronged by being caused to suffer it. Many believe that because P_1 is liable to some harm, he would be wronged to a lesser degree than P_2 by being caused to suffer the additional harm n .⁸ One might argue for this by analogy with a common view about punishment—namely, that it is less bad to inflict a certain punishment on a wrongdoer that is beyond the punishment he deserves than it is to inflict the same punishment on a person who deserves no punishment at all.

Third, suppose there are two groups of people, the A-people and the B-people and that some but not all of the A-people (though it is not known which ones) will be harmed unless a third party takes action to prevent that harm. The third party’s preventive action will unavoidably cause harm as a side effect,

⁷ Joanna Mary Firth and Jonathan Quong, “Necessity, Moral Liability, and Defensive Harm,” *Law and Philosophy* 31 (2012): 687.

⁸ For a related view, see Thomas Hurka, “Liability and Just Cause,” *Ethics and International Affairs* 21 (2007): 199–218.

though this harm will be proportionate because it will be much less than that which the action will prevent. The third party can determine whether this harm will be inflicted on some of the A-people or on some of the B-people. The total amount of harm inflicted would be the same in either case and the amount of harm that each victim would suffer would be the same. If the harm is inflicted on some of the A-people, some of those who will be harmed would not have been harmed at all had the third party not acted. They, it seems, will be wronged by the action. Similarly, if the harm is inflicted on some of the B-people, each victim will be wronged by the action. Assume that none of the A-people or the B-people are liable to be harmed. It seems that if the third party chooses to act in a way that inflicts the unavoidable side-effect harm on the B-people, each victim will be wronged to a greater degree than the different victims among the A-people would have been had the third party chosen to inflict the harm on them instead. This is because each of the A-people was an expected beneficiary of the third party's action whereas none of the B-people was. Each of the A-people could regard the action as being done for her sake even though there is a chance, inevitably realized in some cases, that she will in fact be a victim rather than a beneficiary.

There seems to be nothing in the view that necessity is internal that excludes or is incompatible with the recognition that, in cases of these sorts, and by implication in other cases as well, it is possible for two people to be caused to suffer equivalent harms to which neither is liable and yet for one to be wronged to a greater degree than the other.

4.3 OBJECTION 3: IF NECESSITY IS INTERNAL, THREATENER HAS A RIGHT OF SELF-DEFENSE

A third objection to the view that necessity is internal is that it can imply that, if a victim engages in effective but unnecessary self-defense against a highly culpable threatener, the threatener is permitted to engage in self-defense against her.⁹ This is because any harm that the defender would inflict on the threatener beyond what is necessary is harm to which the threatener is not liable and that would therefore wrong him. Assuming there is no other justification for the infliction of that harm, it seems that the defender makes herself liable to counterdefense if she attempts to inflict it. In *Second Leg*, for example,

⁹ Firth and Quong, "Defensive Harm," 689–90. Helen Frowe, "The Role of Necessity in Liability to Defensive Harm," chapter 8 in this volume, writes that "I share with Firth and Quong the intuition that internalism [that is, the view that necessity is internal] is undermined by its granting to [Threatener] a permission of counterdefense akin to that of an innocent person, should Victim use unnecessary force against him."

the view that necessity is internal seems to imply that Victim makes herself liable to harm that would be a necessary and proportionate means of preventing her from wounding Threatener's second leg (though not to harm that would be a means of preventing her from wounding Threatener's first leg, for Threatener is liable to be wounded in one leg). Critics of the idea that necessity is internal find this implausible. They may also point out that in the original version of *Standing One's Ground*, the difference between the harm that Victim attempts to inflict—death—and the harm that is necessary for her to avoid being killed—none—is the full harm of death. But it seems highly implausible to suppose that Victim makes herself liable to any harm that might be necessary and proportionate as a means of preventing her from killing Threatener, when Threatener is not liable to any of the harm involved in being killed. For that would seem to imply that Victim is liable to be killed in self-defense by Threatener or in other-defense by a third party. And that is intuitively highly implausible. Critics of the idea that necessity is internal seem better able to deal with this issue. They have a plausible explanation of the intuitive view that Victim is not liable to be killed in defense of Threatener—namely, that Threatener has made himself liable to defensive action by threatening to kill Victim and killing him is an effective and proportionate defensive act. And he has no right of defense against a defensive harm to which he is liable.

It does not seem implausible, however, to suppose that Victim makes herself liable to a lesser defensive harm as a means of preventing her from killing Threatener unnecessarily. Suppose, for example, that Victim is about to shoot Threatener when she could safely flee instead. It seems that impartial, disinterested Third Party could permissibly prevent her from killing Threatener by striking the gun from her hand in a way that would injure the hand, provided that Victim could then still flee to safety. And it seems that Victim would not be wronged by being harmed to this extent, given that she would otherwise have acted wrongly by killing a person unnecessarily. She is, rather, liable to have her hand injured as a means of preventing her from wronging Threatener by killing him gratuitously—that is, the injury to her hand would be proportionate in the narrow sense. But if necessity is external to liability, so that Threatener is liable to be killed and would not be wronged by being killed, it is hard to see how it could be permissible for Third Party to harm Victim even in this limited way to prevent her from killing him.

Suppose that the only way that Third Party can prevent Victim from killing Threatener is to kill her. As I noted, it is implausible to suppose that Victim is liable to be killed as a means of preventing her from wrongly killing Threatener. Third Party must therefore not kill Victim but must allow her to kill Threatener. For Third Party, harm is unavoidable: either he will kill Victim or Victim will kill Threatener. Threatener is more responsible for the fact that

harm is unavoidable than Victim is; for although Victim is acting wrongly, Threatener is acting even more wrongly. Given that the unavoidable harm is also indivisible and irreducible—that is, it must all go to one or the other—it ought, as a matter of justice, to go to Threatener. (If the harm could be divided between them, with much the greater share going to Threatener, they might both be liable to a certain share. Or, if the harm could be greatly lessened by harming Victim rather than Threatener—for example, by striking and injuring her hand—then Victim might be liable to that harm and Threatener liable to none at all.)

These claims suggest a problem for the understanding of liability that I have defended. I claimed above that in *Standing One's Ground*, Threatener is not liable to be killed. But I have also just claimed that in the choice that Third Party faces between killing Victim and allowing Victim to kill Threatener, the unavoidable harm ought, as a matter of justice, to go to Threatener rather than Victim because he bears greater responsibility than she does for the fact that harm is unavoidable in the circumstances. Yet I have also claimed in earlier work that in situations in which unavoidable harm is indivisible and irreducible and more than one person bears responsibility for this fact, the one who bears the greatest responsibility is liable to suffer the harm.¹⁰ These claims, however, involve a contradiction: for I claim both that Threatener is not liable to be killed and that he satisfies this sufficient condition for being liable to be killed in this case.

One might say that Threatener is initially not liable to be killed because Victim can safely flee. At that point all harm is avoidable. Thus Threatener is not made liable simply by wrongly threatening to kill Victim. It is only when Victim decides to stand her ground and kill Threatener and Third Party sees this and has the option of killing her that Threatener becomes liable to be killed; for only then does harm become unavoidable, so that the one who is most responsible for this becomes liable. But this is absurd. It implies that because Threatener is initially not liable, Victim would wrong him by killing him; but once she decides to kill him, thereby making some harm unavoidable, she thereby makes Threatener liable to be killed, so that she would not wrong him by killing him.

The solution to this problem lies in the fact that in some cases, a person can be liable vis-à-vis some agents but not liable vis-à-vis others. In *Standing One's Ground*, Threatener is not liable to be killed by Victim. For Victim, harm is entirely avoidable, for she can flee to safety without harming Threatener. Because it is unnecessary for her to kill Threatener, he is not liable to be killed

¹⁰ See, for example, Jeff McMahan, "Who is Morally Liable to Be Killed in War," *Analysis* 71 (2011): 544–59.

by her and will be wronged by her if she does kill him. Third Party's situation is quite different. For him, harm is unavoidable, indivisible, and irreducible. If he does not kill Victim, she will kill Threatener. Given that Threatener bears greater responsibility for this situation than Victim does, Threatener is morally liable to be allowed by Third Party to be killed by Victim. That is, while Threatener is not liable to be killed by Victim and will be wronged by her if she kills him, he is liable not to be saved by Third Party and will not be wronged by Third Party if the latter allows him to be killed rather than killing Victim in his defense.

Similar claims apply to self-defense by Threatener. Although he will be wronged if he is killed by Victim (since she has the option of avoiding all harm), he is, because of his greater responsibility, morally liable to allow himself to be killed by her rather than kill her. This is just to say that because of his greater responsibility for the fact that harm is unavoidable from his perspective, he lacks a right to kill Victim in self-defense. It is not that he altogether lacks a right of self-defense. Suppose that he could prevent Victim from killing him by striking the gun from her hand, thereby unavoidably injuring her hand, but that in doing so he would foreseeably injure himself in a way that would incapacitate him, so that he could no longer threaten her. It seems clear that he could permissibly defend himself in this way. Injuring her hand would be a proportionate response to her wrongdoing, even given his own greater wrongdoing. She would, in other words, be liable to suffer that lesser harm as a means of preventing her from wrongly killing Threatener. But killing her would be a disproportionate means of preventing her from killing him, given his greater responsibility.

This fact illustrates a point that I have sought to defend in earlier work: namely, that there is a comparative dimension to liability.¹¹ If Threatener did not bear greater responsibility for the threat he faces than Victim herself does, Victim would be liable to be killed as a means of preventing her from wrongly killing him. It is only because he bears greater responsibility that she is not liable to be killed and he thus lacks a right to kill her in self-defense.

That it is their comparative degrees of responsibility that make it impermissible for Threatener to kill Victim even though she will otherwise wrongly kill him is supported by a variant of *Standing One's Ground* in which Threatener is not culpable but only minimally responsible for the threat he poses to Victim. Suppose, for example, that Threatener initially has reason, though not decisive reason, to believe that unless he kills Victim, she will wrongly kill an innocent person. He decides that the balance of reasons

¹¹ Ibid.

favors killing her but he is mistaken: she poses no threat to anyone. Although he is not culpable, he has chosen to risk being mistaken in his judgment about Victim. Suppose further that Victim, knowing about Threatener's mistake and knowing also that she has the option of a safe retreat, nevertheless chooses to stand her ground and kill Threatener. Third Party knows all the facts about both and can either allow Victim to kill Threatener or kill Victim in Threatener's defense. Because Victim is acting culpably whereas Threatener is not, it seems that Victim bears greater responsibility for the fact that harm is unavoidable for Third Party. It therefore seems that she is liable to be killed as a means of preventing her from killing Threatener. In short, whether Victim is liable to be killed or Threatener is liable to be allowed to be killed depends on which one is more responsible for the fact that Third Party must choose between them.

Before turning to other objections to the idea that necessity is internal, return briefly to the earlier claim that it would be permissible for Threatener to injure Victim's hand as a means of self-defense. If that is indeed true, it is in part because that option would greatly reduce the harm when some harm is unavoidable from Threatener's perspective. The reduction in harm is sufficient to outweigh his greater responsibility, given that Victim also bears significant responsibility for the unavoidability of harm from his perspective. But there are other ways in which the harm might be reduced by Threatener's harming Victim rather than allowing her to kill him. Suppose, for example, that Threatener is twenty years old while Victim is ninety-five and will certainly die within a few months in any case. Assuming that the harm she would suffer in being killed would be only the loss of a few months of life while his expected loss might be more than seventy years, it might not be disproportionate for him to kill her in self-defense, even given his greater responsibility for the threat he faces from her.¹²

4.4 OBJECTION 4: IF NECESSITY IS INTERNAL, THREATENER IS OWED COMPENSATION

A final objection to the idea that necessity is internal is that if a culpable threatener is wronged by being harmed effectively but unnecessarily, it seems that he is owed compensation by his intended victim.¹³ This implication is

¹² This will be rejected by those who think that the proportionality of defensive harm should be assessed ~~not~~ by comparison not with the amount of harm averted but with the magnitude of the *wrong* prevented. See the discussion of this view below in section 6.

¹³ See Frowe, chapter 8.

thought by some to be absurd. But just as it is not absurd to suppose that a culpable threatener has a limited right of self-defense against wholly unnecessary defensive harm, so it is not absurd to suppose that a culpable threatener has a limited right to compensation if he is harmed wholly gratuitously—that is, when the harm he threatened could have been avoided without harm or cost to anyone. But just as a culpable threatener's right of self-defense is significantly restricted by his greater responsibility for the threat he faces, so his right to compensation is restricted for the same reason. When a culpable threatener has been harmed in effective but unnecessary defense, he is not entitled to full compensation—compensation that would make him as well off as he would have been had he not been unnecessarily harmed. Both he and the potential victim who has harmed him unnecessarily have acted wrongly. They thus share responsibility for the harm he has suffered and it is just that some of that harm should be transferred to the one who caused it unnecessarily. But because the culpable threatener's responsibility is greater, his share of the harm should be greater as well, and most of the harm he has suffered should remain with him.

5 A Possible Qualification

One might think that the reason that necessity is internal is the same as the reason that effectiveness is internal: namely, that both ineffective and unnecessary harming are gratuitous and that gratuitous harming always wrongs its victim. This does indeed seem right if the traditional interpretation of the necessity condition is correct. According to that interpretation, unnecessary harm is always gratuitous, even if it is effective, for the same defensive end could be achieved with an equal or greater probability (and degree of completeness) in a way that would cause less harm, or no harm, to the threatener.

But harm that is unnecessary according to the trade-off interpretation of the necessity condition is not always gratuitous. There are of course cases in which the implications of the traditional interpretation and the trade-off interpretation coincide, so that harm that is unnecessary according to both interpretations is also gratuitous. *Standing One's Ground* is such a case. But there are also cases in which harm that is necessary according to the traditional interpretation is unnecessary according to trade-off interpretation. In these cases, the harm that is unnecessary according to the trade-off interpretation is not gratuitous. An example of such a case is the following:

Paraplegia

As a result of negligence that is only minimally culpable, Threatener will cause Victim to become paraplegic unless Victim acts in self-defense. She has two defensive options. She can either (1) kill Threatener, which will give her a 100 percent probability of avoiding paraplegia, or (2) act in a way that will not harm Threatener at all and will give her a 99.9 percent probability of avoiding paraplegia.

Assume that both options are proportionate. According to the traditional interpretation of the necessity condition, option 1 is necessary because option 2 has a lower probability of success. But suppose that option 2 is morally better all things considered. On this assumption, option 1—killing Threatener—is unnecessary according to the trade-off interpretation. But although killing Threatener is morally unnecessary, it is nonetheless causally necessary for Victim to have a 100 percent rather than a 99.9 percent probability of avoiding paraplegia. To kill Threatener would therefore not be to harm him gratuitously. And given that killing Threatener would not harm him gratuitously, it may be that it would not wrong him, even though it is unnecessary according to the trade-off interpretation. It might be, in other words, that Threatener is liable to be killed even though killing him is unnecessary. If so, this is a case in which necessity is not internal to liability.

This is not to say that it would be permissible for Victim to kill Threatener in *Paraplegia*. Those who reject the idea that necessity is internal nevertheless recognize necessity as a condition of permissibility. If they were to accept the trade-off rather than the traditional interpretation of necessity, they would say that killing Threatener in *Paraplegia* is wrong because unnecessary even though Threatener is liable to be killed and thus would not be wronged by being killed.

Even if it is right that necessity is not internal in a case such as *Paraplegia*, it could still be true that necessity is internal in cases, such as *Standing One's Ground* and *Second Leg*, in which harm that is unnecessary according to both interpretations is also gratuitous. It might only be in cases in which harm that is morally unnecessary according to the trade-off interpretation is causally necessary for some good effect that necessity would not be internal to liability, so that a person could be liable to unnecessary harm.

I continue to doubt, however, whether a person can ever be liable to harm that is morally unnecessary, even when it is not gratuitous. It seems to me more plausible to suppose that even when harm to a threatener is unnecessary only on the trade-off interpretation and not on the traditional interpretation, so that it is not gratuitous, it still wrongs the threatener. In cases in which the

trade-off interpretation implies that the additional harm to the threatener in one option is unnecessary even though it is not gratuitous, the reason the harm is deemed unnecessary is that it is *not morally justified* by the reduction in harm to the potential victim in the other outcome. In *Paraplegia*, for example, killing Threatener is not justified by the fact that it would give Victim a 0.1 percent increase in the probability of avoiding paraplegia. This judgment takes into account that he is responsible through his negligence for the fact that harm or some risk of harm is unavoidable.

Here we might, as before in section 2, convert the question to one of proportionality by ignoring Victim's option of submission and treating ~~her being paralyzed by Threatener~~ as the default option. We can then ask whether her killing Threatener would be proportionate as a means of preventing him from imposing on her a 0.1 percent probability of becoming paraplegic, given that his action involves only minimally culpable negligence. If the answer is that killing him would be disproportionate, that would explain why, of the two defensive options in *Paraplegia*, the second is morally better. And this explanation suggests that if Victim were to choose to kill Threatener, she would thereby wrong him.

6 Proportionality Is Internal to Liability

The reason why the judgment that killing Threatener would be disproportionate suggests that killing him would wrong him is that proportionality is itself internal to liability. This is, I think, more immediately apparent than that necessity is internal to liability. Yet there are cases in which it may seem that proportionality is not internal.

Two Murderers

First Murderer is about to kill Victim. She has a gun but only one bullet. If she uses it to kill First Murderer, she will immediately be killed by Second Murderer.

Given that if Victim were to kill First Murderer she would inevitably be killed a moment later by the independent action of Second Murderer, the only life of which she would be deprived by First Murderer is the short interval before she would have been killed by Second Murderer. If First Murderer were to kill her, he would not deprive her of a long and happy life because that is not what she would have if he were not to kill her. Assume that if Victim kills First Murderer, the brief interval before she is killed by Second Murderer will be worth living,

though less valuable than an average moment in an ordinary person's life. She will therefore gain only a tiny benefit from killing First Murderer.

I suspect that most people would believe that First Murderer is liable to be killed, so that Victim would not wrong him by killing him. Yet it seems that it would be disproportionate for her to kill him. Although killing him would be effective in preserving the moment of life worth living of which First Murderer would otherwise deprive her, it seems disproportionate to deprive even a highly culpable threatener of an entire life as a means of preserving only a moment of good life. This becomes apparent when we consider the following:

Culpable Pincher

Pincher maliciously intends to kill Victim. He reasonably believes that she is a hemophiliac and that if he gives her a hard pinch on the arm, she will die an agonizing death. But in fact she does not have the disease and he has in any case been misled about the vulnerabilities of hemophiliacs. If he pinches Victim, she will merely experience a moment's sharp pain. The only way in which Victim can prevent Pincher from pinching her is to kill him.

Assume that Pincher is even more culpable than First Murderer and that Victim knows about his false beliefs. She knows that he intends to murder her but will fail to do more than cause her a moment's pain. I think it is clear that she may not kill him in self-defense. But assume that having a moment's sharp pain is worse than the loss of a moment of life that, while worth living, would not be particularly valuable. In that case, killing First Murderer must be disproportionate if killing Pincher would be; for the latter is more culpable than the former and threatens a greater harm.

It is compatible with these claims that both First Murderer and Pincher are liable to be killed. Proportionality could be a constraint on permissibility that is external to liability. But we should think more carefully about Pincher. Although he intends to kill Victim, he does not in fact threaten her life. If he is not killed, he will not deprive her of a long future life that would have been worth living but will cause her only a moment of pain. Even though he is highly culpable, it is hard to believe that he has forfeited his right not to be killed as a means of preventing him from pinching Victim. It is more plausible to suppose that he is liable only to some much lesser harm and would be wronged by being killed (though the lesser harm to which he is liable might be greater than that to which an ordinary malicious pincher would be liable).

Now consider First Murderer. He does threaten Victim's life. But in the circumstances the *harm* he would inflict on her in killing her would be less than that which Pincher would inflict on his Victim (given the stipulated

assumption that a moment of sharp pain is worse than the loss of a moment of ordinary experience). If killing Pincher would wrong him by inflicting disproportionate harm on him, it seems that the same should be true of killing First Murderer. And if killing Pincher would wrong him despite his extreme culpability (again, stipulated to be greater than that of First Murderer), it seems that the same should be true of killing First Murderer.

The obvious response to this is that, unless he is prevented by defensive action, First Murderer will *kill* Victim. This is not true of Pincher, and that is the important difference between them. Thus, Helen Frowe, in her contribution to this volume, writes of a similar case in which a potential victim can kill only two of a larger group of murderers that he

might not prevent other people from killing him by killing two of his attackers, but he prevents *those men* from killing him. There's no reason to think that the targets are not liable to be killed and that such killings could not be all-things-considered permissible on most accounts of self-defense. Such accounts don't typically demand that one's defense be necessary for making oneself better off overall, but only necessary for averting a threat. I may, for example, kill a (culpable) person who will otherwise kill me even if I'm certain to die of an illness the next day.¹⁴

One way to try to defend the claims that Frowe advances here is to insist that proportionality is not a relation between the harm that a defender causes and the harm she seeks to avert but is instead a relation between the harm that a defender causes and the gravity of the *wrong* she seeks to avert, or the strength of the *right* she seeks to protect. Frances Kamm seems to presuppose this understanding of proportionality when she writes, in another context, that "a response to multiple wrongdoers can satisfy narrow proportionality so long as the response of each is proportional to his wrongdoing. . . . One compares the wrong to be avoided with what would have to be done to each wrongdoer one at a time."¹⁵

One might think that this conception of proportionality would be congenial to me in particular, as I have argued at length in earlier work that the gravity of the wrong involved in killing does not always, or even usually, vary with the degree to which the victim is harmed by being killed.¹⁶ Yet the gravity of

¹⁴ Frowe, chapter 8.

¹⁵ Frances M. Kamm, *Ethics for Enemies: Terror, Torture, and War* (Oxford: Oxford University Press, 2011), 133–34.

¹⁶ Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (New York: Oxford University Press, 2002), chap. 3.

the wrong involved in an act of killing may not be what is relevant in determining how much harm it is proportionate for a potential victim to inflict on a potential killer in self-defense. It might be, for example, that although it would be gravely wrong to murder a person a few hours before she would unavoidably die of other causes, it would nevertheless be disproportionate and therefore wrong to kill the murderer, depriving him of a great many years of life, just to prevent the loss of those few hours.

The cases of First Murderer and Pincher offer a way of thinking about this problem. I have acknowledged that First Murderer threatens to violate Victim's right not to be killed. It seems that Pincher threatens to violate only her right not to be pinched. Each can be prevented from violating Victim's right only by being killed. If the assessment of proportionality requires weighing the harm caused to the threatener against the strength of the right that is threatened (or the gravity of the wrong to be averted), it seems that killing First Murderer could be proportionate but that killing Pincher could not be. Yet Pincher, like First Murderer, is attempting to commit murder. He is even more culpable than First Murderer and would, in the absence of defensive action, actually inflict a slightly greater harm on Victim than First Murderer would. The difference between them is that Pincher is deluded about his chosen means of killing whereas First Murderer is not. It seems implausible to suppose that Pincher would be wronged by being killed but First Murderer would not.

Alternatively, one might argue that the wrong that would be averted by killing Pincher is in fact not just a pinch but attempted murder, and that the right he threatens is not just the right not to be pinched but also the right not to be the object of attempted murder. These claims highlight the similarities between Pincher and First Murderer. But if they are right and attempted murder is indeed a grave wrong, the conclusion would seem to be that it would be proportionate to kill Pincher, which seems highly implausible. I think we should instead conclude that because in both cases killing the threatener would prevent only tiny harms, killing would in both cases be disproportionate and would wrong those who would be killed.

Perhaps the best way to show that proportionality is internal to liability is via a *reductio* of the contrary view. Consider an ordinary malicious pincher. He dislikes Victim and wants to cause her a bit of pain by giving her a sharp pinch. By approaching her in order to pinch her, the pincher makes himself liable to defensive harm. But suppose the only way Victim can defend herself is to kill him. It seems clear that killing him would be disproportionate and that she must therefore submit to being pinched. But suppose she kills him instead. She has acted wrongly. But if proportionality is not internal to liability, she has not wronged him by the infliction of disproportionate harm. On this view, by threatening to pinch her, he made himself liable to be harmed in

her defense—not to be harmed to a certain degree but simply to be harmed. He has forfeited his right not to be harmed as a means of preventing him from pinching her. If killing him is a necessary means, it would not violate any right he has. This, however, seems clearly mistaken. By merely threatening to pinch Victim, this person does not forfeit his right not to be killed.

My conclusion is therefore that all three constraints on defensive harming I have considered—effectiveness, necessity, and proportionality—are internal to liability.¹⁷

¹⁷ I am very grateful to Nicola Kemp and Jonathan Quong for valuable comments on an earlier draft of this paper.