

Proportionality in Defensive Harm*

Jonathan Quong

April 12, 2013

Albert is angry with you because you won't go out on a date with him, and now he's advancing toward you with very clear intent to do you physical harm. Under these conditions, Albert has made himself *liable* to defensive harm. He has forfeited some of the rights he normally possesses against having harm imposed on him by others, and so if someone imposes harm on him to protect you from his attack, he is not wronged by the imposition of this harm.

Now suppose the only thing Albert is going to do is break your finger. In order to prevent him from breaking your finger, you kill him. Albert, I hope you'll agree, was not liable to *this* level of defensive harm. There may have been a variety of lesser harms you could have imposed on Albert to defend your finger, but killing him is clearly far too severe. Put another way, killing Albert in order to defend your finger is not proportionate. This example illustrates that a person is never liable to defensive harm generally; rather, a person can only be liable to some particular proportionate level of

* For comments and discussion, I am very grateful to Luc Bovens, Susanne Burri, Sarah Buss, Thomas Hurka, Seth Lazar, Christian List, Jeff McMahan, Liam Murphy, Mike Otsuka, Tom Porter, Arthur Ripstein, Jacob Ross, Samuel Scheffler, Nic Southwood, Zofia Stemplowska, and Rebecca Stone.

defensive harm.¹ As David Rodin says, 'For a person to be liable to a harm, just is for that harm to be narrowly proportionate in the circumstances. Proportionality and liability, far from being independent factors, are two manifestations of the same underlying normative relations'.² Proportionality, in the sense that will be the focus of this paper, thus refers to the degree or amount of defensive harm to which a person is liable. Proportionality judgements are an essential part of the individual morality of defensive harm, of the criminal law, and the moral and legal principles that govern just conduct in war. Even if we know *who* is liable to attack in self-defence, or in defence of others, we also need to know *how much* harm it is proportionate to impose in order to assess the moral permissibility of different courses of action.

I have two main aims in this paper. First, I'll explain and criticize a very widely accepted view about proportionality in defensive harm. On this view, the amount of harm it is proportionate to impose on an attacker should depend on the extent to which the attacker is *morally responsible* for creating a situation where unjust harm is unavoidable. The greater the attacker's degree of moral responsibility, the greater the degree of defensive harm it can be proportionate to impose. This idea is, in some form, endorsed by many, probably most, philosophers currently working in this area, but I'm going to argue that we should reject it.

¹ For expressions of this view, see for example Jeff McMahan, *Killing in War* (Oxford: Clarendon Press, 2009), 10; or David Rodin, 'Justifying Harm,' *Ethics* 122 (2011), 79.

² Rodin, 'Justifying Harm,' 79.

The second aim is to offer a positive account of proportionality: *the stringency account*. On this view, the independent variable that explains what counts as proportionate defensive harm is the stringency of the right that the attacker threatens to violate. The more stringent the right that is threatened, the greater the degree of defensive harm that is proportionate. I believe the stringency account provides the best explanation of the essential features of proportionality, and explains our judgements in paradigm cases. It explains, for example, why it would be disproportionate for you to kill Albert in order to prevent him from breaking your finger: the harm you impose on him far exceeds the stringency of the right of yours that he threatens to violate.

Before proceeding to the main arguments, let me clarify a few things. First, I follow Jeff McMahan's general definition of liability to defensive harm.³ On this view, a person is liable to some defensive harm when he forfeits his rights against that harm being imposed, and thus is not wronged or has no standing to complain if the harm is imposed on him. With this general definition in place, we then face the question of what exactly a person must do to forfeit his rights, and thereby become liable to defensive harm. There are competing answers to this question in the literature, but I will rely on

³ See, for example, Jeff McMahan, 'The Basis of Moral Liability to Defensive Harm,' *Philosophical Issues* 15 (2005), 386.

examples where all plausible conceptions will agree that the attacker has acted in a way that entails he is liable to *some* level of defensive harm.⁴

Second, I will focus only on the issue of proportionality as it applies to those who are at least potentially liable to defensive harm. I will thus not consider how we might make judgements about how much harm it is acceptable to impose when defending oneself or others from nonliable people, that is, people who pose a threat despite not having done anything that constitutes a forfeiture of rights. For example, a Nonresponsible Threat who has been blown by an unexpected gust of wind and rendered into a lethal projectile that will land on you is not, I believe, liable to defensive harm. I believe it is still sometimes permissible to impose defensive harm in such cases, that is, to impose harm on nonliable people, but I believe the moral principles governing these cases are sufficiently different that they need to be addressed separately, and so I have nothing to say about such cases here.⁵

Third, proportionality in defensive harm differs in important ways from proportionality in punishment, and so it's important to keep these domains distinct. For example, some people believe that proportionate punishment is *deserved*, but moral desert plays no role defensive harm as I understand it.

⁴ My account of the specific conditions for liability to defensive harm is developed in Jonathan Quong, 'Liability to Defensive Harm,' *Philosophy & Public Affairs* 40 (2012): 45-77.

⁵ I address these cases in Jonathan Quong, 'Killing in Self-Defense,' *Ethics* 119 (2009): 507-537.

Fourth, it is important not to confuse liability with all things considered moral permissibility. Whether someone is liable to defensive harm is one of the important factors that we must take into account when determining whether it is morally permissible to harm that person, but it is neither necessary nor sufficient in deciding whether imposing the harm is permissible all things considered. There are cases when it is permissible to impose harm on nonliable people (e.g. foreseen but unintended casualties arising as part of a justified attack in a just war), and there are cases where it is impermissible to impose harm on someone liable to bear the harm (e.g. when doing so would have sufficiently bad consequences for third parties).

Fifth, most accounts of liability to defensive harm include a necessity condition, that is, they hold an attacker can only be liable to some particular act of defensive harm if imposing that harm is necessary to avert the threat that the attacker poses. I believe that making necessity internal to the notion of liability in this way is a mistake, but I will not argue for, or assume that view here, and nothing in the arguments that follow turn on this issue.⁶ In all the examples to follow, readers can assume the necessity condition has been met, unless I say otherwise.

Sixth, I'll assume that none of the parties are mistaken about any of the facts, and there's no uncertainty regarding what will occur if people perform various actions.

Obviously these are significant simplifying assumptions, but I think it's helpful to begin

⁶ For my view of the necessity condition see Joanna Mary Firth and Jonathan Quong, 'Necessity, Moral Liability, and Defensive Harm,' *Law and Philosophy* 31 (2012): 673-701.

by isolating certain important variables, before looking at additional problems created by ignorance or uncertainty.

Finally, McMahan usefully distinguishes between questions of narrow and wide proportionality.⁷ Narrow proportionality refers to whether some level of defensive harm imposed on a potentially liable attacker (i.e. someone who has acted in a way resulting in the forfeiture of rights against defensive harm) is (or is not) proportionate. Whether it is proportionate for Betty to kill Albert in order to avert his wrongful attack on her is a question of narrow proportionality. Wide proportionality asks whether some act of harm imposition is proportionate once we take into consideration the effects on third parties and other nonliable people. So, for example, whether it would be proportionate for Betty to kill Albert to avert his wrongful attack on her once we also take into account that doing so will foreseeably cause some harm to an innocent bystander is a question of wide proportionality. This paper is exclusively focused on the issue of narrow proportionality.

I. FROM EQUIVALENCE TO RESPONSIBILITY

Before considering the role moral responsibility might play in an account of proportionality, it's helpful to understand the context of the debate—to see how the interest in moral responsibility grew out of earlier, simpler accounts of proportionality.

⁷ McMahan, *Killing in War*, 20-21.

So let's start with a very simple view about proportionality: an eye for an eye, or rather, no more than an eye for an eye. Put a little more generally, this view might give us:

The Equivalence Principle: Proportionality permits the defensive agent (or a third party) to impose no more harm on an attacker than the attacker threatens to impose on the defensive agent.

Although the equivalence principle has the attractions of being simple and somewhat intuitive, it's no longer favoured by philosophers who work on the topic for several reasons. Most obviously, the equivalence principle seems too restrictive in certain cases. Many people believe that it would not be disproportionate for a victim to kill an attacker in self-defence even if the attacker is threatening something less harmful than death. It seems proportionate, for example, to kill an attacker, if necessary, to avoid being violently raped or to avoid being rendered a quadriplegic. Equivalence, at least when measured by the degree of harm imposed, does not seem plausible in many cases.

The equivalence principle also seems insufficiently sensitive to the plurality of considerations that are relevant to proportionality judgements. The equivalence principle focuses only on the threatened *harm* to the victim. But this doesn't seem to be the only thing that matters. Compare the following two cases:

Stroll: Albert is out for a stroll, lost in thought. He negligently fails to notice the sign warning him not to walk any further, since there is a danger of creating rockslides that might injure those below. Albert is about to step forward into this dangerous area, and this will dislodge some rocks that will fall on and break Betty's leg. Carl happens upon the scene. He sees what will occur if Albert steps forward, and he also realizes the only way to stop him in time is to tackle him to the ground, though this will result in some injuries to Albert.

Revenge: Albert is angry that Betty has refused to go on a date with him. He decides to take his revenge by waiting until Betty is walking along the footpath beneath his house, at which point Albert will push some rocks down the hill towards her, breaking her leg. Carl happens upon the scene, and realizes the only way to stop Albert in time is to tackle him to the ground, though this will result in some injuries to Albert.

Although the harm that Betty might suffer in each case is the same—a broken leg—it seems very plausible that our proportionality judgments in these two cases should differ. The amount of harm it might be proportionate for Carl to impose on Albert in *Stroll* should be lower than the amount it would be proportionate for him to impose in *Revenge*. Several differences between the two cases, or the conjunction of

these differences, might explain why. In *Stroll*, Albert is unaware that he poses any threat to Betty, whereas in *Revenge* Albert is very much *aware* of the threat he poses. In *Stroll* the harm Albert threatens to impose is not intended, whereas in *Revenge* Albert's threat is *intentional*. Albert's action in *Revenge* also seems to be a much *graver wrongdoing* in *Revenge* than in *Stroll*. And as a result of these considerations, Albert also seems much *more culpable* in *Revenge* than he does in *Stroll*. It's not clear whether all of these considerations, or only some of them, bear on proportionality, but what does seem clear is that the correct account ought to reflect the fact that these two cases are different. There is more to proportionality judgements than the degree of harm the victim faces.

Recent discussions of proportionality have therefore abandoned the simplicity of the equivalence principle, and instead attempt to identify the full range of considerations that bear on judgements of proportionality in defensive harm. In a recent paper, for example, David Rodin identifies fourteen different variables that he believes are relevant to proportionality.⁸ In one sense this shift from the simple equivalence principle to more complex accounts is a step in the right direction: the best account of proportionality should reflect the underlying complexity of the issue. But it's also true that we want more from an account of proportionality than a list of the different moral

⁸ Rodin, 'Justifying Harm,' 80-81.

variables that may be relevant. I believe a successful account of proportionality should, ideally, meet all of the following conditions:

1. It should provide a plausible explanation of our intuitive judgements about proportionality in paradigm cases.
2. It should be sensitive to the way multiple considerations bear on proportionality judgements, and it should not misrepresent or distort the way these different considerations matter for proportionality.
3. It should offer a coherent framework that unifies these different considerations; it should explain why *these* considerations, rather than others, belong together in an account of proportionality in defensive harm.
4. It should explain the relationship between (a) the necessary and sufficient conditions for liability to defensive harm, and (b) the considerations that determine how much harm a person is liable to bear.

The equivalence principle clearly fails to meet the first two of these conditions, and the more recent and more complex analyses of proportionality, in an effort to meet the second condition, seem unable to meet the third and fourth conditions.

But here is where the appeal to moral responsibility enters the picture. Jeff McMahan offers an account of liability to defensive harm that looks as if it might meet

most, if not all, of the four conditions listed above if we extend it to cover the issue of proportionality.⁹ This is the *moral responsibility account*.¹⁰ McMahan argues that in order for a person to be liable to defensive harm, that person must have voluntarily acted in a way that she could reasonably foresee might result in a threat of impermissible harm to an innocent person. Moral responsibility, on McMahan's account, is thus a pretty thin concept. It doesn't necessarily mean being the subject of praise, blame, or other reactive attitudes. Even if my choice is blameless—for example, my choice to drive to work carefully and conscientiously—I am still responsible for things that might foreseeably occur as a result of that choice (for instance, the brakes might fail and I might get into an accident).

At the heart of the moral responsibility account is an appeal to distributive fairness. If Albert engages in some action or activity that he knows or can reasonably foresee might result in harm to innocent others, and if that risk of harm eventuates and there is now harm to be distributed, it is only fair that Albert should be the person who bears that harm. After all, why should anyone other than Albert have to bear the harm that Albert's own decisions have created? Here we find the familiar and powerful luck

⁹ See McMahan, 'The Basis of Moral Liability,' 394-404; McMahan, *Killing in War*, 157-58, 175-177; or McMahan, 'Who is Morally Liable to be Killed in War,' *Analysis* 71 (2011), 548. A version of this account was first proposed by Kai Draper in 'Fairness and Self-Defense,' *Social Theory and Practice* 19 (1993), 73-92, see especially 84. It is also proposed by Michael Otsuka in 'Killing the Innocent in Self-Defense,' *Philosophy & Public Affairs* 23 (1994), 90-91.

¹⁰ The material in this paragraph and the next is taken from my 'Liability to Defensive Harm'.

egalitarian intuition that individuals ought to bear the costs for their own choices, but should not be held liable for the costs of brute luck, or the responsible choices of others.

McMahan offers this as an account of the necessary and sufficient conditions for a person to be liable to *any* degree of defensive harm. But given the tight connection between liability and proportionality, it's natural to see the moral responsibility account as offering us a way to understand proportionality as well as liability. Suppose Albert poses a threat to Betty, and now it is unavoidable that someone will suffer some level of harm, but this could be either Albert or Betty. Following McMahan and others, let's call this a forced choice situation. There is a cost or burden that must be distributed: what is the fairest way to distribute the burden? An attractive and plausible proposal might be this: the costs should be distributed in a manner that tracks individual responsibility for the creation of those costs. To the extent that one person is wholly responsible for the creation of the costs, that person should bear all those costs. But some of our choices are made under conditions of limited knowledge, or partial duress, or panic, or diminished mental capacity, and thus it would be inaccurate to say we are wholly responsible for those choices in the way we would be if the choices were made under conditions of full information, no duress, etc. Responsibility, in this sense, comes in degrees. When individuals are only partially responsible for the creation of costs, they should (ideally) only be liable to bear the proportion of the costs for which they are responsible. When

no one is responsible for the creation of costs, those costs should be treated as bad brute luck to be shared fairly between all parties.

The variable we use to identify *who* is liable to defensive harm can thus also be the key consideration in explaining *how much* harm the person is liable to bear. Defensive harms are proportionate to the extent that they track an individual's moral responsibility for the creation of a forced choice situation where harm must unavoidably be imposed on someone. Notice that this account can explain the intuitive difference between the two cases I mentioned earlier: Stroll and Revenge. In Revenge it looks like Albert is fully responsible for the threat to Betty's leg, whereas in Stroll, he only bears partial responsibility since the threat to Betty is not something that is connected to Albert's agency in the same strong way as it is in Revenge. There may be many variables that bear on proportionality judgments, but on the view proposed here, they can be at least partially unified and explained by reference to the single concept of moral responsibility. Whether the agent intended the harm, whether the agent was distracted or coerced, whether the agent was mistaken about key facts, whether he is someone with diminished mental capacity — all these issues are united by the way they bear on the *degree* of an agent's moral responsibility for his actions.¹¹ This view of proportionality is captured by the following principle:

¹¹ Some people prefer to use the language of culpability to describe the dimension along which a person's actions can vary depending on the nature and extent of the excusing conditions, but I will generally follow McMahan and use the term moral responsibility or simply 'responsibility'.

The Responsibility Principle: The degree of defensive harm it is proportionate to impose on someone ought to reflect the degree of that person's moral responsibility for creating the forced choice situation that makes it unavoidable that someone must suffer harm.

There are several ways of interpreting this principle, and I'm going to consider both a strong and a weak version of the principle. But the responsibility principle, in one form or another, is widely endorsed by those working on defensive harm and just war.¹² I believe, however, the principle is false. The degree of defensive harm that it is proportionate to impose on a person does not directly reflect that person's degree of moral responsibility for the creation of the forced choice situation. Although moral responsibility is a necessary condition for liability to defensive harm, it plays no direct role in judgements of proportionality.

¹² The following is a partial list of those who suggest that a person's degree of moral responsibility or culpability for creating a situation where harm is unavoidable ought to bear on the issue of proportionality: Kai Draper, 'Defense,' *Philosophical Studies* 145 (2009), 81; F. M. Kamm, 'Failures of Just War Theory: Terror, Harm, and Justice,' *Ethics* 114 (2004), 676-77; Jeff McMahan, 'Self-Defense and Culpability,' *Law and Philosophy* 24 (2005), 766; McMahan, 'Who Is Morally Liable to be Killed in War,' 548; David Rodin, 'Justifying Harm,' 80-84; Victor Tadros, *The Ends of Harm: The Moral Foundations of the Criminal Law* (Oxford: Oxford University Press, 2011), 332; and Suzanne Uniacke, 'Proportionality and Self-Defense,' *Law and Philosophy* 30 (2011), 265.

II. THE STRONG INTERPRETATION

The strong interpretation states that the responsibility principle is the primary basis for all our proportionality judgements in defensive harm cases, that is, whenever we judge that some degree of defensive harm is proportionate, the responsibility principle should be the main grounds for this judgement. To my knowledge, no one has explicitly argued for the strong interpretation of the principle. But it is worth considering since it has some *prima facie* plausibility—especially if you accept moral responsibility as the criterion for liability to defensive harm—and it would, if successful, meet the four conditions for an account of proportionality listed in the previous section.

Let's consider how the strong version might apply to the simple sort of case described at the outset: Albert is responsible for threatening to wrongfully break your finger, and in order to avert his threat, you kill him. Clearly this constitutes a disproportionate level of defensive harm. Can the responsibility principle, on its own, explain this conclusion? The proponent of the responsibility principle might argue that in wrongfully threatening to break your finger, Albert foreseeably creates a situation where a certain amount of harm must unavoidably be distributed, an amount of harm roughly equivalent to a broken finger. Let's give this level of harm a numerical score of 10. What Albert foreseeably does, in other words, is create a situation where 10 units of harm must be distributed, and thus it's fair to hold Albert liable to suffer roughly 10 units of harm. But killing Albert, let's suppose, constitutes 1,000 units of harm. Albert,

however, did not foreseeably create a situation where 1,000 units of harm must be distributed, and this is why killing Albert in self-defence is disproportionate.

This might sound plausible, but there's an obvious problem. Suppose that the only way you can stop Albert from wrongfully breaking your finger is by killing him, and suppose this fact is foreseeable to Albert when he launches his wrongful attack—he knew you would need to kill him to save your finger. Under these conditions, since Albert is wholly responsible for creating this forced choice situation, why is it unfair to require him to bear 1,000 units of harm when the alternative will involve his imposing 10 units of harm on you (who bear no responsibility for the forced choice situation)? It looks like the responsibility principle now tells us it's not disproportionate to kill Albert to save your finger. This conclusion, I assume, is unacceptable. But suppose the principle is revised as follows:

The Revised Responsibility Principle (RRP): When some harm is unavoidable, we ought to choose the distribution that minimizes the total amount of harm where the measurement of harm is weighted to reflect responsibility-sensitivity. Harm is weighted to reflect responsibility-sensitivity by discounting harms imposed on those who are responsible for creating the forced choice situation.

I should note that the RRP is my own attempt to make sense of how the responsibility principle might work in its strong form—it's not something that others have explicitly endorsed.

Now consider our case. The RRP directs us to choose the distribution with the lowest amount of harm, weighted to reflect the fact that Albert is responsible for the forced choice situation. Obviously everything turns on how we discount the harms to Albert to reflect his responsibility, but it's plausible to suppose that the correct way to discount harms to Albert will yield the result that killing him is more harmful than his breaking your finger. Even if each unit of harm to Albert is discounted quite radically, say one unit of harm to Albert is equivalent to ten units of harm to you, the RRP will direct us to choose the distribution where you suffer the broken finger as opposed to the distribution where Albert is killed by you: this will be the distribution that minimizes the total (weighted) amount of harm.

It is true that the RRP is no longer a principle purely concerned with responsibility-sensitivity; instead it incorporates considerations of efficiency (in the sense of minimizing harm), but the way harm is measured reflects the value of responsibility. It's not unreasonable to think an account of proportionality should be sensitive to the total amount of harm suffered, so this modification seems well-motivated and not ad hoc, and it allows the RRP to deliver intuitively plausible results in a range of cases.

But the RRP confronts a fatal problem when we consider cases involving multiple attackers. Suppose 1,000 attackers are intentionally threatening to wrongfully kill you, and each is fully responsible for his action. Imagine they each must press their own individual button, and this collective act of coordinated button-pressing will kill you, unless you press your own button, which kills them all. Many people (including me) share the intuition that all the attackers are liable to be killed by you if this is the only way to avert their wrongful threat to your life. Put another way, it would not be disproportionate for you to kill 1,000 fully responsible wrongful attackers to defend yourself from their attack. More strongly, I don't think there's an upper limit to the number of such attackers that might be liable to be killed by you if necessary.¹³

But the RRP cannot explain this result. In order to explain why it is proportionate to kill any number of fully responsible attackers we would have to assume that the harms imposed on each responsible attacker are completely discounted, that is, they count for nothing when we are calculating which action will minimize the total amount of responsibility-discounted harm. If we didn't entirely discount the harms they each

¹³ Perhaps this does not seem the correct conclusion. But it's important not to confuse two separate intuitions we might have about such cases. One intuition is about whether it would be permissible, all things considered, to kill any number of fully responsible attackers should this be necessary to avert their lethal threat. If you have the intuition that there is some number of responsible attackers that it would be impermissible to kill, this intuition need not be grounded in the idea that it is disproportionate in the narrow sense to do so. It might be that each attacker is liable to be killed (he is not wronged nor are his rights violated), but there are *impersonal* reasons not to carry out killings on such a large scale. Harm may still be impersonally bad even when imposed on someone who has no right against its imposition. I speculate that if it would be impermissible to kill a very large number of responsible attackers, it will be for impersonal reasons and not because the attackers are not liable to be killed.

suffer, there would be some number of attackers that would be big enough to make it disproportionate to kill them. If there is no upper limit to the number of such attackers you can kill, it must be because harms to them don't count at all. But this is not consistent with how the RRP works. The RRP assumes that harms imposed on liable attackers *do* count, but are discounted to reflect their moral responsibility. Moreover, this way of counting (but discounting) the harms suffered by liable attackers is necessary to explain our proportionality judgements in cases like the one described at the outset, where the only way to stop Albert breaking your finger is to kill him.

I believe no strong version of the responsibility principle can explain both of the following:

- A. Why it is disproportionate for you to kill Albert to avert his fully responsible wrongful threat to break your finger, even if this is the only way to avert his threat.
- B. Why the proportionate harms imposed on each fully responsible attacker cannot be aggregated together to place an upper limit on the total number of such attackers it might be proportionate to harm in self-defence.

I think these are central judgements about proportionality in defensive harm, and so I think any strong interpretation of the principle cannot succeed. In order to explain the upper limit for proportionality judgements in standard cases like A, any version of the

principle will need to give some weight to the harms imposed on fully responsible attackers in calculating the total amount of harm brought about by an act of self or other defence. But in order to explain why there is no upper limit to the number of fully responsible attackers that it can be proportionate to harm, the harms imposed on such attackers must be entirely discounted in any such calculation. If responsibility does play a role in judgements of proportionality, its role must be more modest.

III. THE WEAK INTERPRETATION

Recall the original formulation of the responsibility principle:

The Responsibility Principle: The degree of defensive harm it is proportionate to impose on someone ought to reflect the degree of that person's moral responsibility for creating the forced choice situation that makes it unavoidable that someone must suffer harm.

On the weak interpretation, this principle represents one important ingredient for judgements about proportionality in defensive harm, but it does not purport to be the main basis of proportionality judgements. For any given case of defensive harm, other considerations—apart from responsibility—set the upper limit as to what would constitute a proportionate level of defensive harm: this is the amount it would be

proportionate to impose on a fully responsible attacker. Once this upper limit has been determined, we then turn to the responsibility principle to determine if this upper limit is applicable given the facts about a particular case. The less responsible the attacker is, the lower the degree of defensive harm it would be proportionate to impose.

The weak interpretation avoids the objection pressed against the strong interpretation. The weak interpretation cannot be faulted for failing to explain why it is proportionate to kill 1,000 responsible attackers in self-defence (or disproportionate to kill Albert to save your finger) because the weak interpretation does not aspire to explain proportionality judgements as applied to fully responsible attackers. But the weak interpretation still faces objections. In what follows I present four problems for the weak interpretation. Taken individually, no particular objection is decisive. But collectively, I think they give us strong reasons to doubt that the responsibility principle, even in its weaker form, is sound.

A.

First, and most obviously, the weak interpretation no longer offers a unified account of proportionality: it only purports to represent one part of the truth about proportionality in defensive harm. If, other things being equal, we have reasons to prefer accounts that can offer a greater degree of coherence and explanatory unity, then we have reasons to prefer an account that doesn't need to appeal to different principles to explain different

proportionality judgements, but can instead offer a unified account of our judgements about proportionality in defensive harm.

B.

Second, to the extent the responsibility principle does unify at least *some* judgements about proportionality, it does so in a way that fails to capture the qualitative differences between different considerations. Consider the following pair of cases:

Mob Boss: A Mob Boss grabs an innocent man from the neighbourhood, and tells him that he must attack a local shopkeeper who has been refusing to pay the Mob's protection fee. The man knows that if he refuses to do what the Boss demands, the Boss will have one of his henchman exact vengeance. The man thus reluctantly proceeds to attack the shopkeeper, and now threatens to cause him significant injuries.

Reckless Driver: A man is recklessly driving a little too fast through a residential neighbourhood. As he turns a corner, he threatens to crash into a pedestrian. The impact will cause significant injuries to the pedestrian, though the driver will be fine.

In both cases, the two people are not fully responsible for the threats they pose. In the first case, the man is not fully responsible for the unjust threat of harm he poses to the shopkeeper on grounds of duress. In the second case the driver is not fully responsible for the unjust threat of harm he poses to the pedestrian because he does not intentionally pose this threat. The responsibility principle would have us believe that these two considerations both bear on a single dimension—moral responsibility—which then partly determines proportionality. The more duress the man faces in Mob Boss, the less harm he is liable to bear. Similarly, the less intentional and less reckless the driver's behaviour, the less harm he is liable to bear. On this view, if we hold all the other variables constant, it ought to be possible to vary the degrees of duress and recklessness in the two cases such that each man's degree of moral responsibility for the respective unjust threats is roughly equal, and so both will be liable to the same degree of defensive harm, whatever that might be.

This strikes me as counterintuitive and mistaken. The difference between the man in Mob Boss and the driver in Reckless Driver is not a difference of degree, it is a difference of kind. The wrongful actions they perform belong in different moral categories: one is threatening to intentionally and unjustly attack an innocent person, albeit at the command of a mob boss, whereas the other has recklessly driven at too high a speed in a residential neighbourhood. Because I believe the type of moral duty each person violates is very different, I do not think the actions of the two men can be

placed on a single dimension—moral responsibility—where we can vary the degree of duress in the first case such as to render the two cases equivalent with regard to proportionality. Yet this is what the moral responsibility principle proposes that we do. Instead I think what matters, for proportionality, is the *type of wrongdoing* the attacker has committed, and not the extent to which the person who acted is responsible or culpable for doing so.

C.

In light of the previous objection, a proponent of the responsibility principle might concede that unjust threats of harm must be sharply differentiated according to the type of wrong committed. On this view, perhaps some types of wrongdoing are sufficiently different in kind that no amount of variation along the dimension of moral responsibility could render two attackers liable to the same amount of defensive harm.¹⁴ But the responsibility principle would still be important, since two attackers who commit the same type of wrong, or wrongs in the same broad class, could be liable to very different degrees of defensive harm depending on their varying degrees of moral responsibility. But this modification does not avoid a third problem. Consider the following pair of cases:

¹⁴ Though note this interpretation of the responsibility principle differs dramatically from the way Jeff McMahan has developed his view of liability and proportionality. On McMahan's view, it can sometimes be proportionate to kill even a perfectly conscientious driver whose car suddenly veers out of control and threatens to kill a pedestrian, since the driver is still morally responsible for imposing a risk of harm on the pedestrian. See McMahan, *Killing in War*, 165, 178-82.

Case 1: Albert wrongfully threatens Betty with a broken arm (30 units of harm), and is fully responsible for doing so. Betty can avert his threat only by imposing 90 units of harm (e.g. two broken arms and a broken leg).

Case 2: Albert wrongfully threatens Betty with a broken arm (30 units of harm), but is only partially responsible due to duress. Betty can avert his threat only by imposing 90 units of harm (e.g. two broken arms and a broken leg).

The responsibility principle apparently tells us that Albert is liable to a lesser degree of harm in Case 2 than Case 1 because he is less than fully responsible in Case 2. But now let's suppose that the maximum amount of harm it would be proportionate to impose if Albert is fully responsible (recall, this upper limit is determined by appeal to other considerations) is 90 units of harm: the amount Betty can proportionately impose in Case 1 is the maximum amount. Since Albert is less than fully responsible in Case 2, he must be liable to less harm—let's say he is only liable to 80 units of harm in Case 2. But Betty's only options are to impose 90 units of harm on Albert or suffer 30 units herself. If we assume there are no other relevant considerations, it looks as if Betty must allow Albert to harm her in Case 2, because imposing 90 units of harm would be disproportionate.

This conclusion, I am assuming, will strike many people as unpalatable. Even if Albert is partially excused in Case 2, he still acts wrongfully, so why should Betty have to suffer the costs of his attack when she is innocent? Of course, if it were feasible for Betty to impose only 80 units on Albert, this is what she ought to do, but this isn't an option. Isn't it perverse to say that Betty must suffer 30 units of harm to which she is not liable so that Albert can avoid suffering 10 units of harm (recall he is liable to 80 of the 90 units she can impose) to which he is not liable, especially when he is responsible for this forced choice situation?

A proponent of the responsibility principle, however, might accommodate this reasonable concern.¹⁵ The proponent of the principle can say this: if it were feasible, the responsibility principle would allocate less of the total harm to Albert in Case 2 than in Case 1, but liability and proportionality judgements are not about what distribution is ideally possible, nor are they judgements about what amount of harm people ideally deserve or merit. Rather, they are judgements about which distribution from the

¹⁵ Here is another way the proponent of the responsibility principle might accommodate the concern. The proponent might say that Betty is permitted to impose the full 90 units of harm on Albert in Case 2 by combining a liability justification with a lesser evil justification. That is, Albert is liable to bear 80 units of harm (the liability justification) and then given the choice between Betty suffering 30 units of nonliable harm or Albert suffering 10 units of nonliable harm, there is a lesser evil justification for imposing the harm on Albert rather than allowing Betty to suffer it. McMahan has, elsewhere, proposed this method of combining liability and lesser evil justifications. I think this method of combining the two justifications will sometimes explain the permissibility of various acts of defensive harm, though I do not think it can be deployed by the proponent of the responsibility principle without yielding certain counterintuitive results. But here I want to concentrate only on the question of whether a proponent of the responsibility principle ought to believe that Albert's liability varies in these two cases, and so I set the combined justification aside.

feasible set most closely approximates what responsibility-sensitivity requires. And indeed, McMahan makes exactly this suggestion:

Suppose, for example, that there is a fixed harm that either you must suffer or I must suffer. You bear most of the responsibility—say, 90 percent—for our predicament, though I bear the remaining 10 percent. Assuming there are no other factors relevant to the just distribution of the harm, the ideally just distribution would be for you to suffer 90 percent of it while I suffered 10 percent. If it were possible to distribute the harm in that way, you would be liable to suffer 90 percent. But since the entire harm must go to one of us, the nearest approximation to the ideal is for it to go to you. In the circumstances, you are liable to suffer the entire harm.¹⁶

This explains how Albert can be liable to the same level of defensive harm in Case 1 and Case 2 despite the difference in his responsibility. In each case imposing 90 units of harm Albert is the closest approximation to what the responsibility principle would ideally direct us to do if we could distribute the harm any way we like, while holding the total amount of harm constant.¹⁷

¹⁶ McMahan, 'Duty, Obedience, Desert, and Proportionality: A Response,' *Ethics* 122 (2011), 155.

¹⁷ Of course a critic might ask why the relevant counterfactual involves holding the total amount of harm constant while allowing us to imagine that harm can be distributed in any manner we like. If we are imagining other feasible distributions, why not also imagine worlds where we can change to the total amount of harm? Although this worry does not seem unreasonable, the proponent of the responsibility

So, the proponent of the principle can stipulate that judgements about proportionality must be sensitive to what is feasible in the manner described above, but is this stipulation defensible?¹⁸ In order to answer this question, we need to decide whether (and to what extent) non-normative facts can determine individuals' liability to bear costs or receive benefits. We might, following G.A. Cohen, claim that fundamental principles of distributive justice are ultimately fact-insensitive, and thus if justice ideally requires, for example, a 90/10 distribution of harm, then the fact that this distribution is not feasible cannot change the normative reality that one person is liable to 90% of the harm and another person is liable to 10% of the harm. On this view, it may be true that what we ought to do, given the constraint of feasibility, is distribute all the harm to the person who ideally should receive 90%, but in doing so we ought to acknowledge that the level of harm we impose is disproportionate, and so the person who suffers 100% of the harm when he is ideally liable to only 90% may be entitled to compensation since he has had more harm imposed upon him than is warranted by the correct principles of distributive justice. Interestingly, McMahan seems drawn to this conclusion. In the passage immediately following the material quoted above he goes on to say:

principle has a plausible response, namely, if we can also vary the total amount of harm in imagining what an ideal distribution would look like, then the ideal will always be 0 units of harm for everyone. Because defensive harm is not part of retributive justice, it's always better if we can avoid anyone having to suffer harm. But if the ideal standard is always 0 units of harm for everyone, it won't be a useful benchmark for considering the distribution of some actual harm.

¹⁸ The objection developed in the following three paragraphs is drawn partly from the work of Joanna Firth, though I develop the objection in support of the opposite conclusion to the one Firth reaches. See Joanna Mary Firth, 'Moral Liability to Defensive Violence: Why Blameworthiness Matters' (unpublished manuscript).

In the circumstances, you are liable to suffer the entire harm. Yet even though there was a liability justification for inflicting the harm on you, there is a residual injustice. You have suffered more than your fair share. If it later becomes possible for me to compensate you for your having suffered the 10 percent that ideally I ought to have suffered, it is plausible to suppose that I owe you that compensation as a matter of corrective justice.¹⁹

But I find McMahan's position here puzzling. I do not see how he can endorse both of the following propositions:

1. A person is liable to defensive harm X when imposing X would neither wrong him nor violate his rights.
2. There can be duties of justice to compensate a person for imposing a level of defensive harm on him to which he was liable.

If Betty does not wrong Albert or violate his rights by imposing 90 units of harm on him in Case 2, then why should she have to compensate him after the fact for the 'residual injustice' of imposing more harm on Albert than he would have ideally been required to bear if the feasible set had been otherwise?

¹⁹ McMahan, 'Duty, Obedience, Desert, and Proportionality,' 155.

There are at least three possible answers McMahan might give. First, he could argue that we can have duties of corrective justice to compensate someone even when our actions have not wronged this person or violated his rights. But this seems implausible and at odds with our normal understanding of corrective justice. Second, he could revise his definition of what it means to be liable to defensive harm such that a person can be liable to defensive harm of level X, even if imposing level X would also wrong the person or violate his rights. But this is a radical revision to the concept of moral liability to defensive harm, and it would make it much more difficult to understand what we are marking or signifying when we say a person is liable to defensive harm. Third, he could abandon the second proposition above: he could abandon the view that there can be any residual injustice when we impose defensive harm on someone who is liable to bear that harm (and thus abandon the view that compensation could ever be owed under these conditions). Though this third solution looks the most plausible, it leaves the responsibility principle open to the objection that it is concerned with what is feasible at the expense of being insufficiently responsibility-sensitive: it allows that two attackers who commit the very same type of wrong, but are differentially responsible for doing so, may nevertheless be liable to exactly the same degree of defensive harm.

D.

Here's one further objection. Proponents of the responsibility principle tend to focus exclusively on the duress or other excusing conditions *of the attacker*, while ignoring the fact that the *defensive agent* also often faces duress or other excusing conditions.

Shouldn't the responsibility principle take this latter fact into account when calibrating how much harm it would be proportionate for the defensive agent to impose? Imagine a dialogue between a proponent of the responsibility principle (P) and someone about to act in self-defence (D):

P: Careful, the person whom you are going to harm in self-defence is not fully responsible for the unjust threat he poses to you: he is under some duress. As a result, he's not liable to as much defensive harm as he otherwise would be.

D: What about me!? I can see my attacker is being coerced himself, and I feel bad for him, but since he's threatening to seriously hurt me, I'm acting under duress as well, so there's no good reason to lower the degree of defensive harm to which my attacker is liable. To do that would be to treat us in an oddly asymmetrical fashion. My attacker's duress would count in his favour, lowering his level of liability, but my duress apparently counts for nothing: it doesn't increase the level of harm I may proportionately impose. I say either the duress we each face

bears on the proportionality calculation (but then they probably cancel each other out), or else we don't take duress into account for either of us, which again results in no change to what would be proportionate.

I'm not sure there's a plausible reply available to the proponent of the responsibility principle.²⁰

But more importantly, it's not plausible to suppose that the degree of a person's duress, or any other excusing condition, plays this sort of direct role in shaping the contours of individual rights. Excusing conditions, by definition, apply to acts that are wrongful: you can only be excused for doing something wrong. But then excuses cannot also define what counts as wrong. Suppose you see someone in a dark alley approach, and having been violently mugged recently, you are easily panicked in such situations. Although the person has done nothing particularly threatening, you panic and shoot him, seriously injuring him, though it turns out he posed no threat to you at all. The fact that you are partially excused for your action due to your panic has no bearing on the right of the man in the alley not to be shot. This is as it should be. Excuses excuse, but

²⁰ A proponent of the responsibility principle might suggest that the defensive agent's duress will already have been taken into account when calibrating the upper limit of what would be the proportionate degree of defensive harm to impose on a fully responsible attacker posing a threat of a given type, and thus to take the defensive agent's duress into account again—when applying the responsibility principle—would be a form of double-counting. I don't think it is plausible, however, to suppose that the upper limits of proportionality judgments are calibrated partly by taking into account the degree of duress a defensive agent will face. One problem with this proposal is that different defensive agents can experience different degrees of duress when facing the same threat, and so how would we non-arbitrarily decide which degree of duress should inform our judgement regarding the upper limit?

they don't grant us rights to impose harm that we previously lacked. For the same reason, I think we should reject the responsibility principle's suggestion that an attacker's excuses can decrease the liberty rights of others, by increasing his immunity to defensive harm.

I don't purport to have offered decisive reasons to reject the responsibility principle, but I think the objections raised thus far are serious enough that they should lead us to question the widespread adoption of this principle in the literature. Instead I want to propose a different way of thinking about proportionality in defensive harm.

IV. THE STRINGENCY ACCOUNT

Suppose Albert is posing a wrongful threat of harm against Betty. What factors should we take into consideration when determining what level of defensive harm Betty could impose on Albert without violating the constraint of proportionality? I believe the answer is this: proportionality is determined exclusively by reference to the stringency of the right of Betty's that Albert is threatening to violate.²¹ Put more generally,

²¹ Suzanne Uniacke emphasizes the severity of the injustice the attacker threatens to commit in her discussion of proportionality, but she does not agree that the stringency of the right should be the exclusive independent variable in determining the proportionality relationship. She believes other considerations matter, and also endorses a very different view of proportionality, one where proportionality is both a matter of the degree of harm to which an attacker may be liable, but is also a constraint on permissible action independent of an attacker's degree of liability. See Uniacke, 'Proportionality and Self-Defense,' 271.

The Stringency Account: If an attacker threatens to violate a right with stringency level X, then the level of defensive harm it is proportionate to impose on the attacker increases the greater the degree of X, and also exceeds X.

The stringency of a right, as I will be using the term, refers to the strength or weight of the right-holder's claim.²² The more stringent a right is, other things being equal, the more difficult it is to justify infringing the right in question.²³ So, one way to calibrate the relative stringency of different rights might be to ask the following question: how good would the consequences need to be before it would be permissible to infringe the right in question? For example, it might be permissible to infringe right R1 only if we could, at a minimum, save one innocent person from having his legs broken. But it might be permissible to infringe right R2 only if we could, at a minimum, save 50 innocent people from being killed.

I think three things determine the stringency of a right. First, the stringency of a right depends on the way in which the right protects or reflects the *moral status* of the right-holder. Consider the rights to vote and run for elected in liberal democratic

²² Note that others in the literature on rights sometimes use the term stringency differently. On an alternative account, for example, a right is more stringent the fewer the possible conditions there are under which it would be permissible to infringe the right.

²³ A violates B's right when he transgresses the right and has no valid justification for doing so. A infringes B's right, on the other hand, when he transgresses the right but there is an all things considered moral justification for doing so. For example, innocent civilians who are killed by a terror bomber have their rights violated, whereas innocent civilians who are killed as a foreseen but unintended consequence of a proportionate attack in just war may only have their rights infringed. I follow McMahan in making the distinction in this way. See McMahan, *Killing in War*, 9-10.

societies. Even if someone would not be materially harmed by having one of these rights violated or revoked, the rights are stringent partly because of the way they reflect our moral status as free and equal participants in the democratic process.

Second, the stringency of a right depends on the *severity of the harm* that will befall the right-holder if the right is infringed or violated. Consider the right I have against you that you refrain from taking my umbrella without my permission, as opposed to the right I have against you that you refrain from amputating my arm without my permission. The latter is more stringent than the former because the harm of losing an arm is much greater than the harm of losing an umbrella.²⁴

The third variable that can affect the stringency of a right is the *mode of agency* by which the rights violation or infringement might occur. Compare the following two well-known examples:

Trolley: There is a runaway trolley whose brakes have failed headed down a track where five people are trapped and will be killed unless the trolley is diverted.

Fortunately, there is a side-track onto which the trolley can be diverted, but there is one person, Albert, trapped on this side-track, and Albert will be killed if the trolley is diverted. Betty, an innocent passer-by, has seen and understood the

²⁴ I don't have the space to defend the view here, but I believe the account of harm used in calculating the severity of rights infringements must be objective and non-perfectionist.

whole situation, and she now stands at the switch, where she can turn the trolley toward Albert if she chooses.

Overpass: There is a runaway trolley whose brakes have failed headed down a track where five people are trapped and will be killed unless the trolley is diverted. Fortunately, before the trolley will hit the five people it must pass under an overpass. Albert and Betty are standing on the overpass. Betty knows that Albert's weight, but not her own, will be sufficient to stop the trolley before it reaches the five people. If she chooses, she can topple Albert off the overpass and on to the tracks below where he will be killed, but his body will stop the trolley and save the five.

Many people share the belief that it is permissible for Betty to kill Albert in Trolley, but that it is impermissible for Betty to kill Albert in Overpass. But in both cases she kills Albert in order to save the same number of people. If the two cases really do differ in terms of permissibility (I share the belief that they do) then it must be that Albert's right not to be killed in Trolley is somehow less stringent than his right not to be killed in Overpass. And the only salient difference between the two cases appears to be the mode of Betty's agency in killing Albert. There are at least two ways in which Betty's agency in the two cases might differ. First, in Overpass it might be true that

Betty *intends* that Albert be hit by the trolley: toppling Albert off the overpass where he will be hit by the trolley is a necessary means to the end that Betty aims to bring about. But she need have no intentions regarding Albert in Trolley; in this case Albert's involvement is a foreseen but unintended consequence of Betty's act of turning the trolley.²⁵ Second, in Overpass Albert's body provides Betty with an opportunity to save the five that she otherwise would lack—it's only by making use of Albert's body that she is able to save the five. But in Trolley Albert's presence does not provide Betty with an opportunity to save the five: she would be able to achieve her objective just as easily if Albert was not present. Following Warren Quinn, we can say that Betty's killing in Overpass would be a form of *opportunistic* harm, but the killing of Albert would not have this feature in Trolley.²⁶ If killing Albert is permissible in Trolley, but not in Overpass, it thus appears that Albert's right not to be killed intentionally and/or opportunistically is more stringent than his right not to be killed unintentionally and/or non-opportunistically.

The idea that the stringency of a right depends in part on the mode of agency by which someone might infringe or violate the right helps to explain many widely shared

²⁵ I am uncertain about the best way to characterize the relevance of intentions to matters of liability and permissibility. In this paper I rely on the idea that agent's intention matters in assessing the mode of agency involved in various acts of harm-imposition, but I don't offer any specific account of what I mean when referring to an agent's intention. My current, uncertain, view is that it is an agent's *possible* intentions, rather than her *actual* intentions, that matter for questions of permissibility and liability, but I can't defend that view here, though I am trying to do so in a paper in progress, 'Intentions, Permissibility, and Defensive Harm'.

²⁶ Warren Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Double Effect,' *Philosophy & Public Affairs* 18 (1989), 344.

moral distinctions: between a strategic bomber and a terror bomber; between an accidental fatality and a homicide; and between killing someone in self-defence as opposed to using a bystander as a shield to protect oneself from an oncoming threat. In all these pairs, even if the harm imposed in each case is equivalent, it's widely believed that the latter act from the pair is, other things equal, morally worse. We can make sense of these judgements, and many others like them, by appealing to the fact that the mode of agency by which the right is infringed or violated differs in each pair of cases.

I believe these three variables—moral status, severity of harm, and mode of agency—exhaust the relevant considerations that determine the stringency of a right. But even if you think I'm mistaken about this, this needn't be fatal for the stringency account, so long as you accept the idea that rights can vary in their stringency, and that there are a plurality of considerations that bear on the stringency of rights.

With this brief sketch of stringency in hand, we can begin to develop a stringency account of proportionality. Let's begin with this question: how can we meaningfully compare the stringency of a *right* against a particular level of *harm*? You can harm someone without violating her rights (e.g. when the success of my business causes the failure of yours), and you can also violate someone's rights without harming her (e.g. I enter your house without your permission, and leave without your knowledge having damaged nothing). Given that they mark distinct moral categories, how can we make sense of the stringency account's claim that the amount of defensive harm that can be

imposed on a wrongful attacker can somehow be calibrated relative to the stringency of the right that is threatened?

The answer is this: the level of defensive harm that can be imposed on the wrongful attacker can be estimated by considering the stringency of the right that protects the attacker against that level or type of harm *if* the attacker was in possession of the right. So, to take a simple example, if Albert threatens to intentionally break Betty's arm, then the equivalent level of defensive harm would be Betty's intentional breaking of Albert's arm. If Albert had a right against his arm being intentionally broken, then Betty would be threatening to do to Albert exactly what he is currently threatening to do to Betty. Of course once we decide that imposing this level of defensive harm would be equivalent in this way, we also know that if Betty does impose this level of defensive harm, she will not be violating or infringing Albert's rights—she will simply be harming Albert. This is true because—according to the formulation offered by the stringency view—she will only be imposing a degree of defensive harm on Albert to which he is liable. The stringency account tells us that when an attacker threatens to violate a right with stringency level X , the level of defensive harm it is proportionate to impose on the attacker *exceeds* X , and since Betty's breaking of Albert's arm does not exceed what he threatens to do to her, he must be liable to this harm.

But why should the level of proportionate defensive harm exceed X? Why, in other words, should we reject the equivalency principle of proportionality (recall from sect. I)? If Albert wrongfully breaks Betty's arm, Albert both harms Betty *and* wrongs her, whereas a proportionate level of defensive harm imposed by Betty on Albert will, by definition, harm Albert without wronging him. Thus, if Betty was only allowed to impose an equivalent level of harm on Albert, this would not be genuine *moral* equivalence, since it would entail Betty could only do something less serious than what Albert threatens to do to her. Genuine moral equivalence requires that Betty be permitted to do something to Albert more serious than what he threatens to do to her, to reflect the fact that her defensive action will not be a wrongdoing. Only by increasing what Betty may proportionately do to Albert beyond what he is threatening to do to her, can we appropriately reflect the moral asymmetry between them. Exactly how much more proportionately permits a defender to impose on a liable attacker is not, I think, something that is amenable to precise analysis.²⁷

Although I have, in the preceding paragraphs, focused on the harm an attacker might suffer, it's important to remember that each variable relevant to determining the stringency of a right (moral status, severity of harm, mode of agency) is also relevant in our measurement of the level of defensive harm imposed. When we ask whether the action taken by the defensive agent is proportionate, we must ask not only how severe

²⁷ What I've said in this paragraph follows Uniacke, 'Proportionality and Self-Defence,' 260-62.

the harm is that she might impose, but we must also ask about the mode of agency by which she will impose the harm, and also the way her act would threaten the moral status of the attacker if the attacker retained all of his rights. For example, just as our assessment of what a wrongful attacker threatens to do depends on the intentions of the attacker—e.g. whether he intends or merely foresees the rights violation—we must also attend to the intentions of the defensive agent. Other things being equal, it will be easier to justify the proportionality of imposing of defensive harm that is merely foreseen rather than intended, and it will be easier to justify the proportionality of imposing forms of defensive harm that are merely eliminative as opposed to opportunistic.

The stringency account seems to me to have a number of important virtues. First, the stringency account can explain paradigm cases, like the one where you disproportionately kill Albert to stop him breaking your finger: the harm you impose on Albert far exceeds the stringency of the right of yours that he threatens to violate.

Second, the account also provides a unified explanation of why certain variables bear on our judgements about proportionality. For any given threat we must consider: (i) the way in which the right protects the *moral status* of the right-holder, (ii) the *severity of the harm* that the right-holder will suffer if her right is violated, and (iii) the *mode of agency* by which the rights violation occurs. Given that the mode of agency covers a number of different moral distinctions (e.g. intended vs. foreseen, doing vs. allowing, etc.) there will be at least six, and sometimes more, variables that bear on our

proportionality judgement in any given case.²⁸ But the list of considerations that bear on proportionality judgements is not merely descriptive or *ad hoc*: we do not simply draw up a list of things that seem relevant to our judgements about proportionality. Rather, the variables are united by the way they explain the proportionality relationship between the seriousness of the injustice the attacker threatens to commit, and the degree to which the attacker has forfeit his own rights. The more serious the injustice the attacker threatens, the greater the degree of his rights forfeiture. This also nicely explains relationship between liability and proportionality. The condition for liability to defensive harm is that a person must be threatening to violate the rights of someone else; and the amount of harm that is proportionate varies in accordance with the stringency of the threatened violation.

The stringency of a right does not, however, directly depend on an attacker's degree of responsibility for the wrongful threat, and so the stringency view rejects the responsibility principle. As I indicated in the discussion of the responsibility principle, what matters is the type of wrongdoing the attacker commits, not his degree of responsibility or culpability for committing that particular type of wrongdoing.²⁹

²⁸ As mentioned earlier, David Rodin identifies fourteen variables that he argues bear on judgements of proportionality. Many of the considerations he identifies—though not causal/temporal proximity and aggravating conditions for culpability—will be included in the stringency account. See Rodin, 'Justifying Harm,' 80-81.

²⁹ Here is another way the stringency view differs from many existing accounts of proportionality in the literature: it is not directly sensitive to an attacker's level of *causal contribution* to a threatened harm. What matters is the stringency of a right an attacker threatens to violate, but this is only loosely correlated with a person's level of causal contribution. For example, each person who participates in a mob that is stoning

Denying that an attacker's degree of moral responsibility is directly relevant to proportionality may seem counterintuitive in certain cases, and it is certainly at odds with the dominant view in the literature on defensive harm. But consider the following points. First, in many cases, degrees of responsibility seem irrelevant to what constitutes a proportionate level of defensive harm. Few people deny, for example, that it can be proportionate to kill a partially responsible wrongful attacker who intentionally poses a lethal threat if this should be necessary to avert the threat. Second, we have already seen the serious objections that confront the responsibility principle in both its strong and weak forms. Third, I believe that examples where it might seem most intuitive to make proportionality sensitive to an attacker's degree of responsibility or culpability can be explained by appeal to other variables that are correlated with responsibility, but not reducible to it.

Recall the pair of examples used earlier: Mob Boss (where a Mob Boss coerces a man into attacking a shopkeeper), and Reckless Driver (where a driver reckless goes a little too fast and now threatens to harm a pedestrian). I believe the man attacking the shopkeeper in Mob Boss is liable to a greater degree of defensive harm than the driver in Reckless Driver. Although it might appear that this judgement must be explained by

an innocent woman to death may have a negligible causal role (indeed each person's act may be causally unnecessary) in bringing about the woman's death, but I would argue that each person is nevertheless guilty of violating the woman's right not to be killed, and thus each person is potentially liable to be killed to defend the woman's life. Conversely, a person who drives carefully and conscientiously, but whose brakes suddenly and unexpectedly fail and now threatens to hit an innocent pedestrian plays a major causal role in the threatened harm to the pedestrian, but I would argue that the driver does not threaten to violate any rights of the pedestrian, and thus is not liable to any defensive harm.

the fact that the two men differ in their relative degrees of responsibility, I believe that what matters is the fact that the two cases differ with regard to the stringency of the right at stake. In Mob Boss, the man is threatening to intentionally harm the shopkeeper, whereas the reckless driver's threat is unintentional. And the stringency of a right depends on the mode of agency by which a wrongful harm might be brought about: other things being equal, our rights against being intentionally harmed are more stringent than our rights against being unintentionally harmed. This explains why the driver is liable to a lesser degree of harm than the man attacking the shopkeeper. I think the same explanation applies in many other cases where it might appear that the degree of an attacker's responsibility is playing an explanatory role in our proportionality judgements: this appearance, I suspect, is always illusory.

Earlier on, I said that a successful account of proportionality should, ideally, meet all of the following conditions:

1. It should provide a plausible explanation of our intuitive judgements about proportionality in paradigm cases.
2. It should be sensitive to the way multiple considerations bear on proportionality judgements, and it should not misrepresent or distort the way these different considerations matter for proportionality.

3. It should offer a coherent framework that unifies these different considerations; it should explain why *these* considerations, rather than others, belong together in an account of proportionality in defensive harm.
4. It should explain the relationship between (a) the necessary and sufficient conditions for liability to defensive harm, and (b) the considerations that determine how much harm a person is liable to bear.

The stringency account succeeds in meeting these conditions, and this is why I think it represents the best available conception of proportionality in defensive harm. The key to understanding both liability and proportionality is the idea that one person has contravened a duty of justice he owes to others. Liability and proportionality judgements should be grounded in an account of this relationship between the wrongdoer and his victim. The more serious the breach of the duty is, the greater the degree of the rights forfeiture. The stringency account provides the right explanation of this relationship. Other views, like the responsibility principle, do not. An attacker's degree of responsibility for his wrongful action is irrelevant to the question of how weighty or important the duty of justice is that the wrongdoer has violated. In a slogan, it doesn't matter how wholeheartedly an attacker threatens to violate someone's rights, what matters is the type of right he threatens to violate.

V. CONCLUSION

In closing, I want to connect this discussion of liability and proportionality in defensive harm to a current debate about just war. The orthodox view of just war makes a sharp distinction between *jus ad bellum* and *jus in bello*: the principles governing when it is just to go to war as opposed to those principles governing just conduct once the war has begun. On the orthodox view, the latter principles are justified and applied without reference to the former principles.³⁰ This has a number of implications, most importantly, it yields what is called the moral equality of combatants thesis. This thesis states that combatants are equally situated with regard to the rules of war, regardless of which side they fight on. Regardless of whether you are a soldier who fights for the Allies or for the Nazis, for example, provided you observe the *in bello* rules governing warfare, you act justly. Both sides are equally constrained (and equally protected) by the *in bello* principles: a just combatant must observe the same rules of proportionality and discrimination in conducting attacks as a combatant who fights for the unjust side.

This orthodox view has recently been challenged by a number of philosophers, who offer what I'll call the revisionist view.³¹ On this view, the killing and harming of individuals that takes place in war should be governed by the same moral principles that regulate harming and killing individuals outside of war. But this has radically

³⁰ The most influential version of the orthodox account is offered in Michael Walzer's *Just and Unjust Wars: A Moral Argument with Historical Illustrations* (New York: Basic Books, 1977).

³¹ The most influential exponent of the revisionist position is Jeff McMahan. See his *Killing in War*. For an exchange between Walzer and McMahan, see *Philosophia* 34 (January 2006).

revisionary implications for just war theory. In a nonwartime case, the Thug who threatens his Victim with an unjust assault is not on a moral par with the Policeman trying to protect the Victim. The fact the Thug is the wrongful aggressor means he forfeits a series of important rights, but the Policeman who comes to rescue the Victim forfeits none. Revisionists want to apply this logic to the morality of war. Just as in domestic scenarios, if your cause is not just, then it is not just to harm and kill in support of that cause, whereas if your cause is just, then you do not forfeit your rights by proportionately fighting and killing in support of that cause. Because almost everything unjust combatants do is in support of an unjust cause, those acts of harming and killing are all wrongful and unjust even if the combatants do scrupulously observe the *in bello* principles during the course of the conflict. The revisionists thus argue that because we cannot separate *in bello* conduct from *ad bellum* considerations, we must reject the moral equality of combatants, and once the moral equality of combatants is rejected, many further revisions are required.

I've argued that our judgements about liability and proportionately should be determined by focusing on the stringency of the right that an attacker threatens to violate. If someone's harmful actions do not constitute a rights violation, then that person isn't liable to any defensive harm at all on my view. But this means that our judgements about liability and proportionality must be deeply embedded within a broader account of what practices and behaviours justice permits, in particular, which

types of individual behaviour are consistent with respecting the rights of others.

Another way of putting this point is to say that the principles of defensive harm—principles of liability, proportionality, and necessity—are heavily *justice-dependent*. By that I mean that both the justification and application of these principles depend heavily on many broader judgments about the nature of justice.

Some proponents of the responsibility principle tend to argue that if I voluntarily perform some action that foreseeably contributes to a risk that you will be attacked in war when doing so would be fact-relative impermissible, and if this threat eventuates, this information alone is sufficient to show that I am liable to defensive harm.³² But because I think the principles of defensive harm are heavily justice-dependent, and because what justice permits and requires depends on many different variables, not all of which track the standard of fact-relative permissibility, I suspect we need to know a lot more about the context of my action. We need to know whether I was giving a command to combatants to advance on an enemy position, or voting in an election, or cooking food for soldiers, or providing life-saving medical treatment for an injured soldier. Each of these actions might have the same probability of contributing to an unjust attack, and I might be equally responsible for performing each action. But I think that whether any one of these actions is just, or constitutes a violation of others' rights, will depend on the broader set of rules that regulate the social practice. Sometimes we

³² For the distinction between fact, evidence, and belief-relative standards of permissibility, see Derek Parfit, *On What Matters: Volume I* (Oxford: Oxford University Press, 2011), 150-151.

can pose threats to others without acting in a way that violates anyone's rights; think of the prudent driver whose car unforeseeably swerves out of control and threatens to harm a pedestrian. Other times the threats we pose do threaten the rights of others, but they vary in terms of the stringency of the right that is threatened; a reckless driver doesn't threaten the same kind of right as a murderer.

I've argued that the best account of the morality of defensive harm must be sensitive to these differences. So, even if the revisionists are right about their central claim—that the very same principles of defensive harm apply in normal domestic contexts and in war—this doesn't necessarily vindicate many of the specific revisionist claims about just conduct in war. Once we recognize how deeply enmeshed the principles of defensive harm are in broader judgments about justice and individual rights, I'm not sure how much of the revisionist view of just war theory survives.