



PATRICK TOMLIN*

DISTRIBUTIVE JUSTICE FOR AGGRESSORS

(Accepted 27 December 2019)

ABSTRACT. The individualist nature of much contemporary just war theory means that we often discuss cases with single attackers. But even if war is best understood in this individualist way, in war combatants often have to make decisions about how to distribute harms among a plurality of aggressors: they must decide whom and how many to harm, and how much to harm them. In this paper, I look at simultaneous multiple aggressor cases in which more than one distribution of harm among aggressors is available. I show how such cases pose deep questions concerning the nature, role, and scope of the necessity principle, and its relationship to both liability and narrow proportionality. I argue that a hitherto unrecognised measure – ‘narrow proportionality shortfall’ – and its distribution is relevant in choosing how to distribute harms across aggressors. I then extend this analysis to show how this may help us with a puzzle concerning sequential attacks.

I. INTRODUCTION

The literature on the ethics of war and self-defence will often employ cases in which one attacker threatens one victim. There is much that we can learn from such cases. But war involves many people attacking many people, and so there are some questions, including those concerning distributing harms among culpable or responsible threats, that do not arise in these single-aggressor cases. Here I investigate cases involving more than one aggressor, focusing on cases of multiple simultaneous threats in which more than one distribution of harm is available. Our central question is this: How should we distribute harm among aggressors?

* I am grateful to the Oxford War Group, the ELAC 2017 conference in Oslo, and Benjamin Matheson in particular, for useful comments and discussion. I am especially grateful to Helen Frowe, Kieran Oberman, Kim Ferzan and an anonymous referee for extensive written comments.

The paper is an exploration of this topic. Several different issues will be investigated, and some new ideas floated. What draws these different issues and ideas together is that I try to show how consideration of this question about distributing harm among aggressors can raise deeper questions about, and help us to clarify our views concerning, the relationships between some key concepts. Most importantly, since it is usually thought to be the necessity principle's job to compare our defensive options, this topic poses some important, and under-investigated, questions about the necessity principle. These include important questions about necessity's relationship to liability and narrow proportionality, and the scope of the necessity principle's application, which is possibly a lot more expansive than has been traditionally thought. Having explored these issues through simultaneous-aggressor cases, I then show how the positions reached potentially have implications for a puzzle generated by sequential threat cases.

The article proceeds as follows. In Section II, I introduce some key concepts. In Section III, I introduce two key claims about liability – that necessity is internal to liability, and that liability is a one-to-one relation. In Section IV, through considering multiple-aggressor cases, I seek to show that at least one of these claims is false. Sections V and VI investigate the scope of the necessity principle. Once we accept some seemingly uncontroversial claims about necessity's role in distributing harms across attackers, it becomes very difficult to cordon off necessity from having radical implications. Section VII concerns the relationship between narrow proportionality and necessity. I argue that necessity ought to take account of something I call 'narrow proportionality shortfall'. This analysis is extended by Section VIII in which I argue that necessity must take account of distributive concerns, including the distribution of narrow proportionality shortfall. Section IX takes some key lessons from simultaneous-aggressor cases and applies them to a puzzle concerning sequential aggressors. Section X concludes.

II. LIABILITY, NARROW PROPORTIONALITY, AND NECESSITY: THE BASIC MODEL

There are, broadly speaking, two categories of people we might subject to defensive harm: those who are innocent (at least with

regard to the present situation) and those who are not. There are, of course, many finer distinctions we could make, and there are, of course, disputes about how and where to draw the line. Nevertheless, the line is an important one – whether we are harming an innocent bystander or a liable party makes a big moral difference. In this paper, I focus on harms to culpable and responsible aggressors: people who pose an unjustified threat, and bear enough responsibility for that threat to be considered potentially liable to harm. Three key concepts will be referred to throughout – *liability*, *narrow proportionality*, and *necessity*. The paper is in part an investigation of the relationships between these concepts, particularly in contexts in which we face a plurality of potentially liable aggressors.

When someone is responsible for an unjustified threat, they are standardly thought to have forfeited some of their rights against harm. This means that they are potentially *liable* to some defensive harm. When we are liable to harm, we have no complaint against it, and are not wronged by it. However, as Jonathan Quong puts it, ‘a person is never liable to defensive harm generally; rather, a person can only be liable to some particular proportionate level of defensive harm.’¹ *Narrow proportionality* concerns the proportionality of harms with a liability justification. It sets a maximum limit to the harm that may be done to someone under that justification, and is based on a variety of factors, including the level of harm threatened and the aggressor’s level of responsibility.² *Narrow proportionality* contrasts with *wide proportionality*, which concerns the proportionality of harms done to the innocent with a lesser-evil justification.³

Some people think that, at least in some contexts, when people forfeit their rights they forfeit their rights all the way up to the maximum set by narrow proportionality. For example, Kit Wellman distinguishes between strong and weak rights forfeiture, and endorses strong rights forfeiture in the context of punishment.⁴ According to this view, a criminal may be punished up to the nar-

¹ Jonathan Quong, ‘Proportionality, Liability, and Defensive Harm,’ *Philosophy & Public Affairs* 43 (2015): 144–173, at pp. 144–145.

² Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), pp. 19–29; David Rodin, ‘Justifying Harm,’ *Ethics* 122 (2011): 74–110. Quong denies that responsibility is a factor (‘Proportionality, Liability, and Defensive Harm’).

³ McMahan, *Killing in War*, pp. 19–29.

⁴ Christopher Heath Wellman, ‘The Rights Forfeiture Theory of Punishment,’ *Ethics* 122 (2012): 371–393, at p. 375.

rowly proportionate maximum, regardless of whether this punishment achieves some further good.

However, most, perhaps all, philosophers working on self-defence and war believe in a principle of *necessity*. This means that they do not endorse strong forfeiture (in Wellman's sense) in those defensive contexts. In their eyes, narrowly proportionate defensive harms can be impermissible, even when they do no harm to anyone but the responsible threat. In order to be permissible, defensive harm must also be *necessary* as part of the defensive action. The necessity principle applies most clearly in this kind of case: Imagine I can only fend off an attacker with 20 units of harm. This is narrowly proportionate, and it is thus permissible. But now imagine that a new option becomes available, and I can also fend off the attacker with only 10 units of harm. In that case, the 20-unit option includes unnecessary harm, and is thus impermissible. Since using the 20-unit option would involve *more harm* to the *same person* to achieve the *same end*, we can say that this is a case of *gratuitous harm*.⁵

Necessity and proportionality, as I understand them here, perform two distinct tasks. Options that would require excessive force *given the threat faced* are ruled out as disproportionate. Proportionality therefore compares our defensive options with allowing the threat to eventuate. Necessity selects from among our proportionate defensive options: Options that would require excessive force *given our other proportionate defensive options* are ruled out as unnecessary. Necessity therefore compares our defensive options with one another.

The principle of necessity is not only a principle that limits liability-justified harms. It also limits lesser-evil-justified harms. While I will focus on the relationship between liability and necessity here, it is worth remembering that our account of necessity must suitably limit harmful actions under both justifications.

While there is widespread agreement on the requirements of necessity in cases of gratuitous harm, like the one above, there are nevertheless philosophical controversies concerning the principle. For example, there are more complex cases in which it is less easy to determine whether necessity applies and what it requires – the necessity principle is potentially a lot more complex than we might

⁵ I am grateful to Kim Ferzan and an anonymous referee here.

think from cases of gratuitous harm, and, as we shall see, its relationship to liability is also a matter of controversy.⁶

III. TWO CLAIMS ABOUT LIABILITY

In this section I want to outline two claims about liability. In short form, they are as follows:

Claim 1: Necessity is internal to liability

Claim 2: Liability is a one-to-one relation

I'll explain each in turn. The first claim, that necessity is internal to liability, is essentially the claim that aggressors are wronged by unnecessary harms, since they are not liable to unnecessary harms. Crucially, this means that what harm one is liable to is in part a matter of what harm is necessary. Others believe that necessity is external to liability. That is, aggressors are liable to narrowly proportionate though unnecessary harms, but it is nevertheless impermissible to impose them. Therefore, while the defensive agent who imposes unnecessary harm acts impermissibly, she does not wrong her aggressor, and the unnecessary harm does not give rise to a complaint.⁷

The second claim is that liability is a 'one-to-one relation.' This is a claim that liability is only concerned with the attacker and their attack. For example, Jeff McMahan states that 'Liability justifications appear to work by "pairwise comparisons" – that is, by considering each choice between harmings [that is, whether the attacker or the victim is harmed] on its own and determining which of the parties the unavoidable harm ought, as a matter of justice, to be imposed. (This is part of the explanation of why narrow proportionality does not take account of numbers).'⁸ David Rodin claims that 'Liability is

⁶ Seth Lazar, 'Necessity in Self-Defense and War,' *Philosophy & Public Affairs* 40 (2012): 3–44; Jeff McMahan, 'The Limits of Self-Defense' in Christian Coons and Michael Weber, eds., *The Ethics of Self-Defense* (New York: Oxford University Press, 2016).

⁷ For overviews of these issues, see: Helen Frowe, *Defensive Killing* (Oxford: Oxford University Press, 2014), ch. 4; Joseph Bowen, 'Necessity and Liability: On an Honour-Based Justification for Defensive Harming,' *Journal of Practical Ethics* 4 (2016): 79–93.

⁸ Jeff McMahan, 'Liability, Proportionality, and the Number of Aggressors,' in Saba Bazargan and Samuel Rickless, eds., *The Ethics of War* (New York: Oxford University Press, 2017).

a localized comparison between persons in a situation of conflict...values outside that relationship are irrelevant.’⁹

Here, then, is a commonly accepted claim about liability – it is a one-to-one relation. However, consideration of multiple aggressor cases should lead us, I will argue, to either reject this claim or to downgrade its importance.

IV. NECESSITY AND LIABILITY

Consider this simultaneous threat case:

Case 1: Victim is attacked by two unconnected persons simultaneously, who both intend to kill her, and will do so unless she defends herself. Both are equally responsible. Repelling either attack will cause the other attacker to abandon their attack. Attacker 1 can be repelled with 10 units of harm. Attacker 2 can be repelled with 20 units of harm. Both amounts of harm would be narrowly proportionate.

Proponents of strong rights-forfeiture would appear to be committed to the idea that it is permissible to harm Attacker 1, Attacker 2, or both. Since both have forfeited their rights, and there is no necessity condition, both are liable to be harmed up to the proportionate maximum. Most will not accept this conclusion. Most will argue that Victim may harm Attacker 1 and only Attacker 1. Why? Because this minimizes harm. Since the instruction to harm Attacker 1 is an instruction to minimize harm from within the proportionate options, it seems to be the necessity condition that requires Victim to harm Attacker 1 and not Attacker 2.

This is *not* a case of gratuitous harm, since the two potential harms are to two different people. It is a case of how to distribute harm between two liable parties. But it seems that necessity nevertheless has something to say (and most would accept that it is necessity that applies here – after all, if necessity had nothing to say whenever harms might befall different persons, it would hardly apply at all in the context of war). Yet, if necessity does apply here, this case tells us something very important about necessity and its relationship to liability.

Recall that it is often thought that necessity is internal to liability (Claim 1), and that liability is a one-to-one relation (Claim 2). But if it

⁹ David Rodin, ‘Justifying Harm,’ p. 99. Rodin doesn’t here use the language of one-to-one relationships, but I think it is a plausible way to read this claim, given his other work, especially on bloodless invasions. However, even if Rodin has a wider definition of ‘situation of conflict’ in mind, we will need an account of what is within and what is outside that situation – something which, as I will show, is tricky to provide.

is *necessity* that tells us that we ought to harm Attacker 1 and not Attacker 2 in Case 1, then at least one of these two claims must be false. We have a trilemma, in that we cannot, I argue, endorse the claim that necessity tells us to harm Attacker 1, and Claim 1, and Claim 2.

In Case 1, if our judgment that Attacker 1 ought to be harmed is based upon necessity, that judgment can only be reached by comparing harms *across attackers*. The harm to Attacker 2 is unnecessary, and that level of harm is unnecessary due to the option of inflicting less harm on *somebody else*. Therefore, necessity is not *simply* about pairwise comparisons. Since necessity is not simply about pairwise comparisons, then if necessity, in part, determines what harm we're liable to (i.e., if Claim 1 is true) then liability is not simply a one-to-one relation (i.e., Claim 2 is false). If, on the other hand, Claim 2 is true, and liability can *only* be established by pairwise comparisons between each Attacker and Victim, then Claim 1 must be rejected – necessity is not internal to liability.

We could claim that Attackers 1 and 2 form a 'unit', and so this is a single attack, not two separate attacks. Therefore, liability is not quite a one-to-one relation, but it is still limited to those involved in the 'situation of conflict'. In order to claim this, a defender of this view will need to give us an account of how to demark 'situations of conflict'. But even if they were able to do so, the conjunction of Claims 1 and 2 would still be in trouble. Consider this case:

Case 2: Attacker attacks Victim threatening to kill her. Victim can defend herself by killing Attacker, or scratching Bystander.

It seems that necessity would require us to harm Bystander. All else equal we might think Victim is required to harm Attacker, or let herself be harmed, over harming Bystander. But all else is not equal here. Either Victim or Attacker will end up dead, *or* Victim can very mildly harm Bystander. Bystander is clearly not part of the 'situation of conflict' (she is, after all, a bystander) and yet necessity seems to be telling us that she must be harmed. So we again see that necessity looks across the various harmful options, singling one out. And again, not all of these harmful options will harm people involved in the attack. If necessity helps to determine liability, we must reject the claim that liability is a one-to-one relationship. If liability is a one-to-one relationship (or involves only those who are part of the situation

of conflict) we must deny that necessity helps to determine liability.

We could keep Claims 1 and 2 by denying that it is *necessity* that tells us who to harm in Case 1 and Case 2. On this picture, necessity is more limited, focusing simply on what is necessary *to stop this attack* (or perhaps it is even more limited, and restricted to ruling out gratuitous harm only). We might then reason through Case 1 as follows. The 20-units harm to Attacker 2 is proportionate and necessary. The 10-units harm to Attacker 1 is also proportionate and necessary. And so both attackers are liable to those amounts of harm, since necessity and liability both, on this picture, concern *this* attacker and the threat that they pose. However, we might think, there is some further principle, *external to liability*, which states that, of the two necessary harms, one is to be chosen over another. This might be a more general ‘minimizing harm’ principle.

One thing to note here is that by postulating two separate principles – necessity, which is internal to liability, and a minimizing harm principle, which adjudicates between harms to which persons are liable – this view proposes what seems to be a radical revision to defensive morality. The job that the minimizing harm principle is supposed to do is one that necessity has standardly been thought to do – i.e., selecting the minimally harmful option from among our proportionate defensive options. Perhaps they are parts of the same principle – in the way that McMahan has carefully and usefully distinguished between narrow and wide proportionality, perhaps there is a difference between narrow and wide necessity. Narrow necessity is internal to liability, while wide necessity adjudicates between harms to which persons are liable (and, perhaps, other harms too).

Here is one potentially attractive entailment of this type of ‘split necessity’ view: since both Attacker 1 and Attacker 2 are liable, it allows us to say that if Victim chooses to harm Attacker 1, she acts impermissibly, but she does not *wrong* Attacker 1. On this view, you wrong someone when you do more harm than is necessary *to them* to *stop their attack*, but you do not wrong them when you fail to minimize harm more generally.

However, this view, as currently stated, cannot save the conjunction of Claims 1 and 2. To see this, consider Case 1 again. The position explored above relies on each of the two Attackers being

liable to harm. But, according to this new view, neither attacker is, in fact, liable to harm. If necessity is internal to liability, and liability is a one-to-one relation, then we must look only at the threat each Attacker poses, and whether the harm *to him* is necessary to prevent it. On this measure, neither harm is narrowly necessary. There is a way to stop Attacker 1 *without harming him at all* – namely, harming Attacker 2 – and there is a way to stop Attacker 2 *without harming him at all* – namely, harming Attacker 1. Therefore, to harm either is to harm *him* more than is necessary.

The view could be amended, however, to fix this problem.¹⁰ If narrow necessity rules out comparison with options that harm third parties, then both options in Case 1 are narrowly necessary. In considering whether it is narrowly necessary to harm Attacker 1, I cannot appeal to the fact that I could harm Attacker 2, since Attacker 2 is a third party insofar as the threat posed by Attacker 1 is concerned.

This reformulated view, however, also has some concerning implications. First, consider a case in which I can either kill my Attacker or run away. Killing the Attacker would be impermissible and would wrong the Attacker, since it would be narrowly unnecessary. But now imagine that my running away would very mildly harm a Bystander. I would still be morally required to run away, but now I would *not* wrong the Attacker by killing him. It seems odd to say that in both cases I am required not to kill the Attacker, but I only wrong him in one case. Second, now consider a case in which I could kill my Attacker, or mildly harm a Bystander, but that harming the Bystander would have enormous benefits (for example, it would save my life *and* five further lives). I would be required to harm the Bystander, but, on this view, would not wrong my Attacker if I chose to kill him, since killing him would be narrowly necessary. Now add an option to run away without harming anyone. According to the view under consideration, the addition of this *impermissible option*, an option I would be morally required to ignore, would make it the case that I now *would* wrong my Attacker if I chose to kill him, since killing him is now not narrowly necessary.

A further general effect of separating necessity into narrow and wide necessity – with wide necessity operating external to liability –

¹⁰ I am grateful to an anonymous referee for comments here.

is that, as with having an external necessity principle more generally, this limits the significance of liability: Showing someone is liable, and that nobody else will be harmed, is not enough to justify harming them.

There are further worries with the idea that narrow necessity is internal to liability and that liability is a one-to-one relation. Consider this case with one attacker and two victims:

Case 3: Attacker attacks Victim 1 and Victim 2. Individually, each attack would license self-defence up to the narrowly proportionate limit of 10 units of harm. Only shooting the Attacker in the legs, imposing 15 units of harm, will prevent Attacker from completing his attacks. Bystander must decide whether to impose 15 units of harm.¹¹

Now, this case may seem unproblematic. Intuitively, we clearly ought to handle the case as follows. First, the narrowly proportionate limit for the *combined* attacks is 20 units. Second, the necessary force is fifteen units. Therefore, the Attacker is liable to 15 units of force.

But recall that we are investigating the conjunction of the claims that necessity is internal to liability, and that liability is a one-to-one relation. The above, intuitive response relies on *combining* the attacks to establish liability. However, if Claims 1 and 2 are correct, then liability for each attack must be established independently (because liability is one-to-one) and we must establish whether the harm necessary to stop *each attack* is narrowly proportionate *given the aim of stopping that attack* (since narrow necessity is internal to liability). So, we first ask whether 15 units is necessary and proportionate for preventing Attack 1. It is necessary, but not proportionate. And so Attacker is not liable to that harm in order to prevent Attack 1. The same is true of Attack 2. And so he is not liable to that harm either. Since he is not liable to 15 units in order to prevent Attack 1, nor 15 units to prevent Attack 2, he is not liable to 15 units in order to prevent both attacks.

To be clear, this counter-intuitive outcome is not generated by assuming that liabilities cannot be aggregated. If, for example, Bystander could prevent the attack by inflicting 10 units that would prevent Attack 1, and, separately, 10 units that would prevent Attack 2, this would be permissible, since both these harms would be

¹¹ I came to realise this problem with the one-to-one conception of necessity through engaging with Kimberly Kessler Ferzan's exploration of liability and desert, 'Defense and Desert: When Reasons Don't Share,' *San Diego Law Review* 55 (2018): 265–290. Ferzan's work in turn engages with Uwe Steinhoff's work.

necessary and proportionate *to preventing the particular attack*, and so Attacker would be liable to both. (For example, if Bystander could shoot Attacker in both arms, independently causing him to drop the guns that threatens each Victim.) The trouble comes because of the conjunction of Claims 1 and 2. If necessity were external to liability (i.e., if Claim 1 were rejected), we could do things in this way: Attacker is liable to 10 units for Attack 1, and 10 units for Attack 2. Therefore, he is liable overall to 20 units of harm. A separate principle, external to questions of liability, namely necessity, confirms that the necessary harm of 15 units falls below the *total* amount of harm to which he is liable.

However, if we reject Claim 2, and instead seek to establish Attacker's *joint liability* for *both* attacks, then Attacker is potentially liable to up to 20 units of harm (combined narrow proportionality) and actually liable to 15 units (the harm necessary to prevent both attacks). Liability is established by joining the two attacks together.

Consider, again, what combining Claims 1 and 2 together seems to require of us.¹² Since we must establish liability on a one-to-one basis, and since we must keep necessity internal to liability, then liability must be established on a one-to-one basis. Thus, we are forced to say that 15 units of harm is too much to impose in order to stop Attack 1 and so is disproportionate. And the same is true of imposing 15 units in order to stop Attack 2. However, two separate harms – of 10 units each – would be allowed if they independently stopped each attack. So, if Claims 1 and 2 are true, then we would be allowed to impose 20 units of harm on Attacker, but not 15. This is an odd result, for surely necessity should require us to inflict 15 units of harm over 20 units of harm *in order to achieve the same end*.

Our consideration of the above cases reveals the following. First, necessity (or something like it) must surely consider various harmful options that spread across attacks. Second, it follows from this that one of the following is true: (1) liability is not a one-to-one relation, (2) that necessity is external to liability, or (3) necessity is not the only principle that helps us to adjudicate between proportionate defensive options. Third, if (2) or (3) is the case, then there are external limits on one-to-one liability which seek to limit harm across attacks. If this is true, the claim that liability is a one-to-one relation is

¹² I am grateful to Helen Frowe for useful discussion here.

less powerful than we initially might be tempted to think. For liability is not, then, enough to grant us permission to harm, even if nobody who is not liable is in the frame.

Although I have speculated here that what we call ‘necessity’ may actually be an amalgamation of concerns or principles, I will largely leave that speculation behind from this point on. Whatever we call it, we are morally required to minimize harms among liable attackers, and this goes beyond the instruction to avoid gratuitous harm. I want to investigate this requirement further, and will call it ‘necessity’ for ease.

V. SEPARATE ATTACKS

It is tempting to see narrow proportionality and necessity as two limits on defensive harm, with the main difference between them being that narrow proportionality is concerned with features of the threat (‘what you did was out of proportion given the threat I posed!’) while necessity is concerned with the comparative merits of the options available to the defensive agent (‘you didn’t need to do *that*, you could have done *this*!’).

We now see, however, that, regardless of whether they are internal or external to liability, narrow proportionality and necessity appear to work quite differently in other ways. Narrow proportionality looks at the one-to-one relation between attacker and victim (though it can be aggregated, as in Case 3), while necessity seeks to minimize harm across narrowly proportionate defensive options, sometimes directing us to harm one attacker rather than another (as in Case 1), or a bystander over an attacker (as in Case 2).

Thus far we have considered cases in which the attacks were in some senses separate, giving rise to different one-to-one relations, but the defensive options were linked by producing the same effects. However, if we extend the analysis thus far, we get some potentially controversial results. Consider Cases 4 and 5.

Case 4: Attacker 1 attacks Victim, threatening 100 units of harm. The attack can be prevented by inflicting 100 units of harm on Attacker 1. Attacker 2 also attacks Victim, also threatening 100 units of harm. The attack can be prevented by inflicting 50 units of harm on Attacker 2. Victim can only prevent one of the attacks.

Case 5: Attacker 1 attacks Victim 1, threatening 100 units of harm. The attack can be prevented by inflicting 100 units of harm on Attacker 1. Attacker 2 attacks Victim 2, also threatening 100 units of harm. The attack can be prevented by inflicting 50 units of harm on Attacker 2. Neither victim can prevent the attack against herself or the other, but Bystander can prevent one.

The analysis thus far suggests that the requirement to minimize harm stretches across attacks, and attackers. Thus, it would seem to instruct Victim in Case 4 to inflict 50 units of harm on Attacker 2. She is going to suffer 100 units of harm either way, and so the requirement to minimize harm states that she must inflict 50 units of harm rather than 100 units of harm in order to avoid an additional 100 units of harm to herself. Some may find this counter-intuitive. They may think both responses are both narrowly proportionate and necessary, and that Victim can choose as she wishes whom to harm. Both attackers have made themselves liable and may be harmed. But the problem with such a view is that it seems to rely on an account of necessity that was rejected in the previous cases. Cases 1–3 show that necessity, or something like it, gives us guidance that we ought to minimize harm *across attackers*.

This expansive understanding of necessity seems even more controversial when the attacks are fully separate, as in Case 5. In this case, our analysis thus far suggests that Bystander ought to save Victim 2, since doing so minimizes harm and all else is equal.

If we transport this expansive view of necessity to war, we will generate some very controversial conclusions. On this view of necessity, if we have two wars we might enter in defence of others, both of which will do the same amount of good, we would be required to enter the war which would do the least harm to culpable aggressors, even though both wars, considered independently, are necessary and proportionate. This is controversial. Especially since there is no principled reason to restrict our options to war. Scholars have noted a controversy about proportionality and what counterfactual is involved: When states decide to go to war, they need to compare 'going to war' with 'not going to war'. How, several philosophers have asked, should we fill out 'not going to war'? It could be a very demanding standard: A war might be proportionate only if the good it achieves *compared with the very best options available to the state* is worth the death and destruction involved.¹³ But

¹³ David Mellow, 'Counterfactuals and the Proportionality Criterion,' *Ethics and International Affairs* 20 (2006): 439–454; Thomas Hurka, 'Proportionality and Necessity' in Larry May and Emily Crookston eds., *War: Essays in Political Philosophy* (Cambridge: Cambridge University Press, 2008); Kieran Oberman, 'War and Poverty,' *Philosophical Studies* (forthcoming; <https://link.springer.com/article/10.1007/s11098-017-1012-4>).

philosophers have been slower to notice that there is also a similar issue related to necessity.¹⁴ Necessity requires us to *compare options*, but how should we construct our option set? On the expansive view currently under discussion, a war might be necessary only if there is no other way to achieve a comparable level of good or prevention of harm without causing as much harm. Therefore, necessity *demand*s that we should, for example, alleviate poverty over going to war.

This may seem overly demanding. We can avoid these demanding entailments if we allow that necessity can apply across persons, but demand that it be restricted to a particular ‘situation’. If we can cordon necessity off in some way, we can say that necessity only applies to the choice between *some* options, and not to, for example, the choice between a war and alleviating poverty. We might then appeal to some independent ‘minimizing harm’ principle that advises us (or requires us) to minimize harm, but is not related to liability. This proposal is attractive. Indeed, some seem to think it is *obviously* the way to handle the cases discussed in this paper. I am attracted to this view, but, as I am going to explain, I think it is hard to get it to work.

VI. LIMITING NECESSITY’S REACH?

Necessity rules on which of our available options we should take. If we are to cordon off necessity, preventing it from applying to a choice between war and alleviating poverty but allowing it to apply to choices about whether to harm an attacker more or less, we must have a principled way of determining the scope of the options to which the necessity principle applies. I think a necessity principle with a limited scope intuitively offers us the best picture, but finding the right way to restrict necessity’s reach is difficult. That is what I want to show in this section. The idea of this section is to lay down a challenge to anyone who denies that necessity dictates that we ought to alleviate poverty rather than go to war: How do we restrict necessity’s reach?

The most obvious way to try to delineate a ‘situation’ or set of options to which necessity applies is to do with the people involved. The clearest examples of the necessity principle being breached in-

¹⁴ A notable exception here is Kieran Oberman, ‘War and Poverty’; and ‘Killing and Rescuing,’ *Philosophical Review* 129 (2020). In these papers Oberman argues for a very expansive view of necessity.

volve harming someone more than we need to. If individuals are wronged by being harmed unnecessarily, we might try to limit necessity's reach by restricting its application to options that harm the same people. For example, we might try this principle:

Limiting Principle 1: Necessity applies only to options which harm the same persons.

This principle is obviously flawed. First, necessity would not apply in this case:

Case 6: Attackers 1 and 2 attack Victim. They can be stopped by killing Attacker 1, or by killing both Attackers.

According to Limiting Principle 1, killing both Attackers would not be condemned by the necessity principle, since the two options available harm different sets of people, and so the necessity principle would not apply across these two options.

Here's another principle:

Limiting Principle 2: Necessity applies only to options in which *some* of the same persons will be harmed.

This principle has four problems. First, necessity would not apply to the decision in Case 1. Second, this principle would render the *jus in bello* necessity condition next-to-toothless. Consider this case:

Case 7: Platoon needs to capture an enemy base. If they attack from the south, they will need to kill one sentry, Sentry 1. If they attack from the north, they will need to kill Sentries 2-21.

This is a clear case of where we would expect the *jus in bello* necessity principle to apply. According to Limiting Principle 2, it does not. So long as killing 20 enemy soldiers is proportionate, then Platoon should feel free to attack from the north or the south. We might try to set aside Case 7 by claiming that since the enemy soldiers all fight for the same army, this is a special case, and that necessity can apply to *groups* of people. But, first, this would fly in the face of much contemporary just war theory. And, second, we would no longer be able to compare harms to our own soldiers and civilians with harms to enemy soldiers and civilians. Third, the necessity condition would not apply to this case:

Case 8: Country A is threatened by Countries B and C. Defending A from either B or C will cause the other to back off. Country D can defeat B with a loss of 100 of B's soldiers, or C with a loss of 10,000 of C's soldiers.

The third problem with Limiting Principle 2 is that necessity should clearly condemn killing when the same objective can be achieved

without killing *anyone*. Limiting Principle 2 only seems to apply to choices between options in which someone is going to be harmed. And if we extend it to options in which nobody will be harmed, the principle will then apply to the choice between killing in war and not killing and alleviating poverty.

The fourth problem with Limiting Principle 2 is that it is open to what I will call ‘overlapping problems’. Since I discuss these in more depth below (concerning overlapping populations of who is to be saved), I will not detail these issues here.

An alternative to focusing on the parties to be harmed is to limit necessity’s reach by focusing on the people to be saved. As Kieran Oberman observes, when we ask whether harm is necessary, it can only be necessary or unnecessary *to achieving some goal*.¹⁵ Since, in cases of both self-defence and other-defence, our goal is to *save people*, we might try limiting necessity by looking at options which save the *same people*. For example, in Case 1, Victim aims to save herself. How much harm, she might ask, is necessary to achieve this aim? Therefore, we might use this principle:

Limiting Principle 3: Necessity applies only to options in which the same people are to be saved.

This principle implies that if we could kill Ted to save 100 people, or kill nobody and save *the same* 100 plus Bill, necessity would not dictate that we choose the latter option, since the same people are not saved in both options. We could amend Limiting Principle 3 to avoid this implication if we allowed necessity to apply only to options in which *at least* the same people are saved (i.e., to options which save some core group and options Pareto superior to those). However, even under this revised reading, in a war context, necessity would have no bearing on this choice.¹⁶

Case 9: Country A aims to save Country B from Country C. This involves sending troops into Country B. A is confident it will win the war, but there will be civilian casualties. If A’s army enters B from the east, 1000 civilians from the eastern district will be killed. If A enters B from the west, 10,000 civilians from the western district will be killed.

Here, since different people will be killed, different people will be saved. If A chooses to enter from the east, the 10,000 westerners who would be killed under the alternative plan will instead be saved.

We could instead try this principle:

¹⁵ Oberman, ‘War and Poverty’; ‘Killing and Rescuing’.

¹⁶ For related discussion, see Oberman, ‘War and Poverty’, section 5.

Limiting Principle 4: Necessity applies only to options in which *some* of the same people are saved.

Here is my worry with this principle. Consider the following ‘overlap case’:

Case 10: Aggressor attacks Victims 1, 2, 3 and 4. Option 1 saves Victims 1 and 2, at a cost of 10 units of harm to Aggressor. Option 2 saves Victims 2 and 3 at a cost of 100 units. Option 3 saves Victims 3 and 4 at a cost of 1000 units.

Now, compared pairwise under Limiting Principle 4, necessity says we must choose Option 1 over Option 2, Option 2 over Option 3, but that we have a free choice between Options 1 and 3, since necessity doesn’t apply to that choice. So, necessity doesn’t necessarily tell us to do anything regarding the choice between Options 1, 2 and 3. Perhaps this is not a problem: Perhaps necessity is not an issue of pairwise comparisons, but of comparing a full range of options. Provided that Option 1 saves some of the same people as Option 2, and Option 2 saves some of the same people as Option 3, perhaps necessity applies to the choice between all three options. However, if we view necessity in that way, then two new worries arise.

The first is that necessity faces a difficulty from the ‘principle of irrelevant alternatives’. According to Limiting Principle 4, provided there is a ‘chain’ of options in which at least some of the same people are saved in adjacent options, then necessity applies across the whole range of options. So, in Case 10, you should select Option 1, since $1 > 2 > 3$. Let’s now remove Option 2. Removing Option 2 *which you weren’t going to choose anyway*, means that you now have a free choice between Options 1 and 3.

More importantly, Limiting Principle 4 does not necessarily mean that necessity would not apply to a choice between war and alleviating poverty. Often a choice between a war and alleviating poverty will involve saving at least one person in common. But even if not, we simply need to be able to construct a ‘chain’ of options (however morally implausible some of them may be) and the options of war and alleviating poverty will then become subject to the necessity constraint. For example, imagine that we are deciding whether to militarily intervene in Country A or alleviate poverty in Country B. Provided we *also* have the (morally implausible) option of devoting all of our resources to saving only Lukas from Country A and only

Max from Country B, then necessity applies, according to Limiting Principle 4, to the choice between war and alleviating poverty.

I don't have any answers here. My intuitions tell me that, at least when it comes to harming the potentially liable, necessity is restricted to particular defensive situations. However, it is very hard to find a principled way of restricting necessity.

VII. NECESSITY AND NARROW PROPORTIONALITY

Necessity and narrow proportionality are usually thought of as independent tests. Narrow proportionality establishes upper limits on harm, necessity tells us to minimize harm. Our job as defensive agents is to check that they coincide: that the option which minimizes harm is below the threshold set by (aggregated) narrow proportionality. In this section, I want to argue that the relationship between narrow proportionality and necessity is potentially more complicated than this.

I will argue that whether a harmful action can be deemed necessary involves looking at something I call 'narrow proportionality shortfall'. If I am right, this shows two things. First, even if necessity is external to liability, the relationships between the various concepts remain complex. We cannot simply say that narrow proportionality sets an internal limit based on pairwise comparisons, whilst necessity sets a fully independent external limit that is concerned with aggregate harm, since narrow proportionality is an important concept in deciding the permissible distribution of harms *across* a plurality of responsible threats. Second, since necessity is sensitive to narrow proportionality, this shows that they are two distinct criteria. This is in contrast to some recent scholarship that has sought to show that necessity and proportionality are but two sides of the same coin.¹⁷

Consider this case:

Case 11: Victim is attacked by two unconnected persons simultaneously. It is narrowly proportionate to impose 200 units of harm on Attacker 1, and 20 units of harm on Attacker 2. Repelling either attack will cause the other attacker to abandon their attack. Attacker 1 can be repelled with 21 units of harm. Attacker 2 can be repelled with 20 units of harm.

In this case, it is more harmful to repel Attacker 1 than Attacker 2.

¹⁷ Lazar, 'Necessity in Self-Defense and War,' pp. 17–18; Thomas Hurka, 'Proportionality in the Morality of War,' *Philosophy & Public Affairs* 33 (2005): 34–66, at pp. 37–38.

So, any view under which necessity merely requires us to minimize harm, or to minimize priority-weighted harm, would instruct Victim to harm Attacker 2. But, in my view, necessity instructs Victim to harm Attacker 1. Seth Lazar and Jeff McMahan argue that necessity does not demand that we minimize harm (all else equal) but rather that we minimize *morally-weighted harm* (all else equal). That is because when we must choose, for example, between a harm to an attacker and a slightly lesser harm to an innocent bystander, we should harm the attacker who has created the situation in which someone must be harmed.¹⁸ My claim here is that harms to culpable aggressors can also have different moral weightings.

Narrow proportionality represents the maximum possible harm that can be justified under a liability justification for achieving some specific good. In Case 11, harming Attacker 1 involves inflicting roughly one-tenth of the harm to which he is potentially liable. Repelling Attacker 2 involves inflicting *all* of the harm to which he is potentially liable. Necessity should take account of this. Just as necessity should distinguish between harms done to the innocent and those done to the liable, so it should distinguish within harms done to the liable. It can do this if the moral weighting of harm takes account of what I call the 'narrow proportionality shortfall' – that is, how far short of the maximum potential liability, or narrowly proportionate limit, a harmful action is. In Case 11, the narrow proportionality shortfall of harming Attacker 1 with 21 units is 179. The narrow proportionality shortfall of harming Attacker 2 with 20 units is zero. This huge discrepancy in narrow proportionality shortfall explains why we should harm Attacker 1, even though this does more harm.

Case 11 shows that necessity is concerned with narrow proportionality shortfall. There are several ways that we might incorporate this concern with narrow proportionality shortfall into a theory of necessity. Here are three. First is the extreme view, in which this concern for narrow proportionality shortfall is the only concern when considering harms to potentially liable parties. In other words, there is no concern with absolute levels of harm, and a concern only with narrow proportionality shortfall. According to this view, narrow proportionality sets a baseline against which to measure harm.

¹⁸ Lazar, 'Necessity in Self-Defense and War,' p. 6; McMahan, 'The Limits of Self-Defense'.

Here is how this view would handle Case 11. Since Attacker 1 is potentially liable to 200 units of harm, the 'baseline' against which we measure harms to him becomes -200 units. Essentially, that means that we will prefer 200 units of harm to him over a single unit of harm to an innocent person (whose baseline remains at 0). Only once we have filled the full quota of harms to which you are potentially liable will we look at harming an innocent person.

Since Attacker 2 is potentially liable to 20 units of harm, his baseline becomes -20 units. Therefore, any distribution of harm in which Attacker 1 has fewer than 180 units of harm is to be preferred to any distribution in which Attacker 2 is harmed at all, because only once Attacker 1 has received 180 units of harm will the two be on equal footing. I think this view has a compelling logic, though it is counter-intuitive. Should we prefer 180 units of harm to Attacker 1 over one unit of harm to Attacker 2?

Second is the proportion view. In Case 11, 21 units of harm is roughly 10 percent of what Attacker 1 is potentially liable to, while 20 units of harm is 100 percent of what Attacker 2 is potentially liable to. We might think necessity dictates that we should prefer whichever harm is a lesser proportion of the harm to which the Attacker is potentially liable. However, this again could generate cases in which we should prefer hundreds of units of harm to one potentially liable attacker over a single unit of harm to another.

Third is the mixed view. On this view, the concerns outlined in either or both of the extreme view and the proportion view are coupled with a concern to minimize absolute levels of harm. I prefer this view. But the central point is that, whichever of these we choose, if we accept that narrow proportionality shortfall matters to necessity, then the following three things follow. First, this is further evidence that, if necessity is internal to liability, then liability is not a one-to-one relation – necessity (and thus liability) chooses which of the potentially liable persons we should harm, based on facts about other persons. Second, it shows that necessity and narrow proportionality are independent concepts, one of which informs the other. Third, it shows that a moral consideration (namely, necessity) prefers harms that are *further away* from narrow proportionality thresholds to those that are close to the maximum liability. If, all else equal, necessity demands that we maximize narrow proportionality short-

fall, then it shows morality to be sensitive to how close harms get to the narrow proportionality threshold. This final point will become important when we consider sequential aggressor cases.

VIII. NECESSITY AND DISTRIBUTION

Thus far we have largely considered cases in which we must choose which of multiple simultaneous aggressors we ought to harm. However, sometimes we might also have choices of harming more than one aggressor, and different distributions of harm may be available. For example, consider this case:

Case 12: Victim is attacked by Attacker 1 and Attacker 2. It will be narrowly proportionate to inflict 100 units of harm on either. Victim has three defensive options: 100 units of harm to Attacker 1; 100 units of harm to Attacker 2; and 51 units of harm to both.

In my view, Victim ought to harm both with 51 units, even though this is more harm overall. This cannot be because we must minimize harm. It also cannot be because of maximizing narrow proportionality shortfall – the total narrow proportionality shortfall of harming one of the attackers is 100, the total narrow proportionality shortfall of harming both is 98.

Most straightforwardly, this case shows that we should be concerned not only with harm, but with the distribution of harm. Before we delve into what this distributive concern is, if we accept that in Case 12 Victim ought to inflict 51 units of harm on both attackers, then we face a conceptual choice. One is to fold this distributive concern (whatever it may be) into necessity – allowing distributive concerns to affect the moral weighting of harm. Another is to introduce such concerns as an independent desideratum.¹⁹ The cost of this position, however, is that necessity is relegated from hard side-constraint to defeasible desideratum, since this independent distributive concern would allow us to depart from necessity's recommendations. Therefore, I think it makes more sense to fold this concern into necessity.

So, some kind of distributive concern should be added to the moral weighting of harm. But this doesn't show that we must be concerned with (say) an equal distribution of harm. This distributive concern might itself be sensitive to narrow proportionality shortfall. Consider the following case:

¹⁹ Lars Christie suggested this to me.

Case 13: Victim is attacked by two unconnected persons simultaneously. It would be narrowly proportionate to harm Attacker 1 with 40 units of harm, and Attacker 2 with 20 units of harm. The following defensive options are available. Option A: Both attackers are harmed with 20 units of harm. Option B: Attacker 1 receives 30 units of harm, and Attacker 2 receives 10 units of harm.

In this case, the total amount of harm (40 units) and the total narrow proportionality shortfall (20 units) remain constant between the two options. All that changes is the distribution of harm (in Option A 20-20, in Option B 30-10), and the *distribution* of the narrow proportionality shortfall (in Option A 20-0, in Option B 10-10). My own view is that necessity requires Option B in this case. Were necessity focused on minimizing narrowly proportionate harm, we'd expect it to be indifferent between Options A and B. If necessity cared about minimizing harm *and* distributing that harm as equally as possible, we'd expect it to prefer Option A.

Here are two views which would account for preferring Option B.

Equality of Narrow Proportionality Shortfall: All else equal, fairness between attackers requires equalizing the distance from the narrowly proportionate limit of harm to which they are potentially liable.

Shifted-Baseline Prioritarianism: Units of harm are not equally bad – they get worse the more harm there is. When someone is potentially liable to some harm, this shifts the 'baseline' against which we measure how bad the harm is. For example, the 201st unit of harm to somebody who is potentially liable to 200 units of harm should be considered of equal badness to the 21st unit of harm to someone who is liable to 20 units of harm, and to the first unit of harm to someone who is not liable to any harm. The first unit of harm to someone who is potentially liable to 20 units of harm equally weighted to the 181st unit to someone who is potentially liable to 200 units.

The idea with Shifted-Baseline Prioritarianism is that harm gets progressively worse, as in standard prioritarianism,²⁰ but narrowly proportionate harms are discounted.²¹ Importantly, though, not all units of narrowly proportionate harm are given the same discount: Those closest to the narrowly proportionate limit are given very little discount. Those further away are given progressively more and more discount. This position would give us the advantage of caring both about total harm, and about total narrow proportionality shortfall (the view that I found most attractive in the previous section) with one underlying explanation – we care about total *morally weighted and priority-weighted* harm, and the priority-weight-

²⁰ Derek Parfit, *Equality or Priority?* (Kansas: University of Kansas, 1991).

²¹ This may, therefore, be a specification of Richard Arneson's 'Responsibility-Catering Prioritarianism.' Richard J. Arneson, 'Luck Egalitarianism and Prioritarianism,' *Ethics* 110 (2000): 339–349.

ing baseline is shifted by narrow proportionality.

I will not investigate further which of these two options we should endorse. It may be that which is preferable turns on deeper questions concerning distributive ethics that I cannot venture into here, although it may be that narrowly proportionate harms are *sui generis*, and we can, say, endorse prioritarianism within the limits of narrow proportionality while rejecting it outside. Similarly, we may endorse equality under normal circumstances, but find the idea of having to make sure we treat attackers fairly unpersuasive.

Instead, I want to draw out several implications of the views suggested. It seems that both counsel that harms closer to the narrow proportionality limit are, in some sense, something to be additionally wary of, or *harder to justify*. They are not the same as harms further from the narrow proportionality limit. In *Equality of Narrow Proportionality Shortfall*, this commitment manifests itself in a concern to fairly distribute the avoidance of these harms, rather than the harms themselves. In *Shifted-Baseline Prioritarianism*, this manifests itself in weighting harms closer to the limit more heavily than those further away.

One implication of this concerns some cases that are, conceptually and normatively, difficult. The most clear cases for the principle of necessity – both in terms of whether it applies as a principle, and in terms of what it recommends – are those in which the exact same good can be achieved by more or less harm to the same person (i.e., what we have called *gratuitous harm* cases). But there are also cases in which our options vary both in terms of how much harm *and* how much good we will do. For example, consider this case:

Case 14: Attacker threatens Victim with 40 units of harm. Victim can prevent 39 units of harm by imposing 60 units of harm on Attacker, or she can prevent all 40 units by imposing 70 units of harm on Attacker. The narrow proportionality limit is 100 units.

In some ways, choosing between these options is clearly an issue for the principle of necessity. Both options are narrowly proportionate, so we have a choice between narrowly proportionate options, and it is necessity's job to select between such options. But in another way, it is more like a proportionality calculation. That is because, in essence, Victim's decision is whether ten units of *additional* harm to

Attacker can be justified to prevent one unit of additional harm to herself.²²

My point here is simply this. If we allow that narrowly proportionate harms are harder to justify the closer to the narrowly proportionate limit they get, this will impact on how we treat such cases: This isn't a case of weighing one unit of harm to an innocent person against ten units of harm to an aggressor, which is how such cases are usually discussed. It is a case of weighing one unit of harm to an innocent (which may itself be priority-weighted) against ten units which range from 40 to 30 in terms of narrow proportionality shortfall. Keeping the unit of harm to the Victim constant, such a unit may justify ten units of harm to an additional aggressor who hasn't yet been harmed, but may *not* justify ten units of *additional* harm to an aggressor who is already going to receive 60 units from a limit of 100.

IX. SEQUENTIAL ATTACKS

Thus far, I have focussed on simultaneous threat cases: cases in which a plurality of aggressors poses threats at the same time. In closing, I'd like to apply some of the lessons that we have drawn from simultaneous-threat cases to sequential-threat cases.

There is an important puzzle concerning sequential-threat cases. According to Frances Kamm, there is no limit to the number of aggressors whom I could kill in order to save my own life. Since killing in order to fend off each attack, taken on its own, is necessary and narrowly proportionate, Kamm reasons that there can be no objection to the collection of defensive actions: 'One compares the wrong to be avoided with what would have to be done to each wrongdoer one at a time, and if there is no violation of proportionality in any individual comparison then there is no violation tout court.'²³

Yet many of us are deeply troubled by the idea that, provided we can kill one aggressor in order to achieve some goal, we can kill an infinite number of aggressors in order to achieve that goal. Is there some way to avoid this implication? There is, of course. Since the implication arises from the picture of self-defence and war that has

²² Lazar, 'Necessity in Self-Defense and War'; McMahan, 'The Limits of Self-Defense'.

²³ F.M. Kamm, *Ethics for Enemies* (New York: Oxford University Press, 2011), pp. 133–134.

emerged from both ancient and recent scholarship on this issue, we could rip up this picture and start again, making sure that the new picture doesn't have the unpalatable implication that we can kill infinite aggressors. This is too hasty – this model serves us well in general, and so we should do all we can to preserve its core conceptual and normative commitments. Ripping it up would be to throw the baby out with the bathwater. Alternatively, we could stipulate that one cannot kill an infinite number of aggressors. But this response is *ad hoc*.²⁴

To forewarn: I do not claim to solve this puzzle in what follows. But I do think that what we have learned about liability and necessity in simultaneous-attack cases might help us with this puzzle.

First, it is important to understand that necessity applies across attacks and is a moral concern for minimizing harm not just *within* a one-to-one relationship but *across* such relationships. Several philosophers have suggested that there may be some countervailing reasons against harming those whom it would be individually narrowly proportionate to harm, but which derive from aggregation of narrowly proportionate harms. For example, David Rodin argues that wide proportionality is concerned with the way that harms to the liable add up, while Jeff McMahan proposes an additional form of proportionality – proportionality in the aggregate.²⁵ McMahan's proposal in particular may seem somewhat mysterious. In defence of the existence of proportionality in the aggregate, McMahan is unable to point toward much other than that it gets the intuitively right answers – in other words, it sails a little too close to the *ad hoc* response to the puzzle.

What we have discovered in the present paper is perhaps partial evidence for the existence of something like proportionality in the aggregate – there are other elements of defensive morality that work in similar ways. They constrain harms and dictate choices of harms between multiple aggressors, even when each harm, were it the only option, would be proportionate, necessary, and permissible. Necessity reveals a concern for the comparative lessening of such harms. But wouldn't it be odd if morality were keen to limit such harm

²⁴ I am grateful to Victor Tadros for discussion which helped me to see these twin dangers.

²⁵ McMahan, 'Liability, Proportionality, and the Number of Aggressors'; Rodin, 'Justifying Harm'.

comparatively but not absolutely? If so, proportionality in the aggregate would seem to be on a somewhat firmer basis.

Relatedly, McMahan in particular takes the apparent one-to-one nature of liability to be what generates the puzzle. But we now see that either liability is not one-to-one (if necessity is internal to liability), or that liability is subject to external checks (if necessity is external to liability). Therefore, there are other parts of defensive morality that already operate in ways that are similar to proportionality in the aggregate. This makes McMahan's response to Kamm's puzzle less mysterious.

Second, in addition to these general points about necessity, we have found a more direct link between necessity and narrow proportionality. Necessity takes account of narrow proportionality. It is concerned with maximizing narrow proportionality shortfall, and with its distribution. Here, I think there is potentially the beginnings of a justification and understanding of the way in which narrowly proportionate actions can add up to a disproportionate course of action. Necessity seems to prefer harms further away from the narrow proportionality threshold than those nearer to it. This suggests that harms close to the narrow proportionality threshold can be justified, but are worse, harder to justify, or to be avoided more, than those further away.

My tentative view is that these observations can provide a firmer backing for the kinds of restrictions on harming responsible threats articulated by McMahan and Rodin. McMahan's proportionality in the aggregate could, I think, be understood as the way in which narrow proportionality shortfall can aggregate. We have seen that necessity prefers harms that are a greater distance below the narrow proportionality line, even if those just under the line could be justified in a single attacker case. If a series of harms are close to the line, then the same concerns that motivate necessity's preference for harms further away from the line could motivate a concern with the aggregation of harms close to the line. The thought here is simple: Necessity is concerned with minimizing 'only just' proportionate harms. If there are differentiations within narrowly proportionate harms, such that we are concerned with how close to the line harms get, we might also be concerned with how these harms aggregate.

In particular, the *Shifted-Baseline Priority View* would seem to lessen the gap between narrow and wide proportionality. It doesn't eliminate the difference, but it does mean that the two kinds of proportionality are not as different as may first appear. What narrow proportionality does, on this view, is (differentially) discount harms to the potentially liable. These harms still count though – they're not morally free shots – and so many harms that are right on the narrow proportionality threshold can aggregate just as harms to the innocent can aggregate. This view of narrow proportionality as a progressive discounting mechanism might therefore provide some backing for Rodin's view.

This also would explain why proportionality in the aggregate can differentiate between culpable and responsible threats. Responsible and culpable threats are both liable to be killed in order to save a life. However, we should differentiate between *de facto* limits to liability and *de jure* limits to liability. *De facto*, the limit will often be the same – killing is just about the worst thing you can do to a person, and both are liable to be killed. But culpable aggressors can be liable to far greater harms than those they threaten to impose on others – for example, it is proportionate to kill a culpable aggressor to avoid being paralyzed, even if one will have many good years as a paralyzed person. Therefore, the *de jure* limit to avoid being killed by a culpable aggressor is far higher than death, and far higher than the *de jure* limit on a responsible threat. Therefore, when we kill many responsible threats, we impart many harms very close to the narrow proportionality threshold. This is not the case when we kill many culpable aggressors (if the good at stake is a person's life). It may be that there is some limit to the number of culpable aggressors we can kill to save a single life (I think there is) but it will be far higher than the number of responsible threats.

I won't pretend that what I have said in this section is any more than a sketch of a defence of McMahan's and Rodin's responses to Kamm's view. But thinking through simultaneous-aggressor cases, and what we learn about liability, necessity, and narrow proportionality from such cases, certainly makes Kamm's logic seem less compelling. In particular, I would now reject the thought that liability is simply a one-to-one relation, and Kamm's view seem compelling in part because of that assumption.

X. CONCLUDING THOUGHTS

This paper has covered a lot of ground. The common thread throughout has been thinking about cases involving more than one attacker, and how we should distribute harms between those attackers. Consideration of these cases raise major issues in terms of how to understand the principle of necessity, its relationship to liability and narrow proportionality, and its scope. In particular, I have argued for the following:

1. Necessity ought to tell us which of two potentially liable parties to harm. It follows from this that (a) necessity is *either* not internal to liability, *or* liability is not a one-to-one relation; and (b) if necessity is external to liability, then necessity, as an external principle, takes account of the distribution of harms across attackers.
2. Necessity possibly has different components or concerns, or is possibly an amalgam of distinct principles (for example, avoiding gratuitous harm, minimizing morally-weighted harm, distributing morally-weighted harm). Some of these concerns or principles may be internal to liability (narrow necessity), some external (wide necessity).
3. Once we realise that necessity applies across individuals, it becomes very difficult to cordon off necessity so that it applies to the choice between some options and not others. We must either supply a principled way of restricting necessity's scope, or accept that it has a far greater reach than traditionally supposed, not only directing us in how and whether we go to war to pursue some end, but whether we pursue that end or another.
4. Necessity should be sensitive to 'narrow proportionality shortfall' – that is, how closely the various harms one is deciding between are to the narrow proportionality threshold.
5. Necessity should take account of distributive concerns, including the distribution of narrow proportionality shortfall. I articulated two principles for doing so, *Equality of Narrow Proportionality Shortfall* and *Shifted-Baseline Prioritarianism*.
6. These findings may help us in thinking about sequential attacks. That (a) liability is not a one-to-one relation, or morality allows external checks on liable harms that seek to limit them, and (b) necessity cares about the distribution of harm and narrow proportionality shortfall, may provide some support for the idea of 'proportionality in the

aggregate' in sequential-attacker cases, a notion which suggests that you cannot kill an infinite number of liable aggressors if your cause would allow you to kill one.

OPEN ACCESS

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

*Department of Philosophy
University of Warwick, Coventry, UK
E-mail: patrick.tomlin@warwick.ac.uk*

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