‘Defensive Escalations’

Gerald Lang¹

Received: 22 March 2021 / Accepted: 3 September 2021 / Published online: 22 September 2021
© The Author(s) 2021

Abstract

Defence cases with an escalatory structure, in which the levels of violence between aggressor and defender start out as minor and then become major, even lethal, raise sharp problems for defence theory, and for our understanding of the conditions of defence: proportionality, necessity, and imminence. It is argued here that defenders are not morally required to withdraw from participation in these cases, and that defensive escalations do not offend against any of the conditions of defence, on an adequate understanding of them. No plausible interpretation of proportionality or necessity excludes defensive permissions in escalatory cases. Moreover, the structure of escalatory defensive cases also sheds useful light on permissible responses to conditional deadly threats. This is because we can analyse conditional deadly threats as disguised cases of escalation. Because defensive resistance is permissible in escalatory cases, it is also permissible in cases involving conditional deadly threats.

Keywords Self-defence · Escalation · Conditional threats · Proportionality · Necessity · Imminence

On the margins of defence theory, we find cases with an escalatory structure. In these cases, the violence on both sides is minor at first, but will escalate unless at least one side disengages, leading to much larger losses later in the sequence. These cases are intuitively puzzling, and it is not obvious what we should say about them.

I have two main aims in this article. First, I want to make clear what defensive permissions are available in such cases, and why. Second, I want to demonstrate that my analysis of escalatory cases can be fruitfully applied to lethal conditional threats, which appear on the surface to have a non-escalatory structure.

In pursuing both these aims, the conditions of defence will be in play. The general presumption that defensive violence is justified in at least some circumstances is carefully regulated by defence theory. Defensive violence is usually thought to

✉ Gerald Lang
g.r.lang@leeds.ac.uk

¹ School of Philosophy, Religion, and History of Science, University of Leeds, Woodhouse Lane, Leeds, West Yorkshire LS2 9JT, UK

be governed by at least three conditions. These conditions—imminence, necessity, and proportionality—are distinct but complementary ways of managing the harms inflicted by defensive action. Defensive violence inflicts harm and injury, after all, and these are presumptive evils which we should aim to avoid if we can. These debates about the conditions of defence are complex, and some of the complexities will be addressed as we go along. For the moment, I shall introduce only simple and relatively unrefined versions of them, so that we have clear targets from which to proceed.

According to the imminence condition, defensive violence is justified only if the threat of violence is imminent. Imminence is better understood as an epistemic condition rather than a strictly temporal one. This is because the moral interest of temporal imminence is ultimately epistemic: the temporal imminence of the threat is typically correlated with the quality of the evidence for it. If Daisy has unimpeachable evidence that Tom will definitely attack her at some point in the future, the temporal distance between now and then does not matter. She can permissibly take defensive action now, so long as the necessity and proportionality conditions are satisfied.

According to the necessity condition, if defensive violence is unnecessary, then it is impermissible. If defenders can evade the threats posed to them by aggressors without violence, then this is what they should do. Defensive violence is permissible only when it is unavoidable, or when the defender cannot evade aggression by retreating from it.

The proportionality condition insists that defensive violence be proportionate, in the sense that the harms inflicted on the aggressor through defensive action are not excessive compared to the harms that the aggressor would impose on the defender in the absence of defensive violence.

Necessity and proportionality can come apart: a proposed course of defensive action might be necessary without being proportionate, and vice versa. Imagine, in *Pinch*, that Daisy cannot avert Tom's attempt to pinch her without killing him. In *Pinch*, lethal defensive violence is necessary, but disproportionate. Tom has no right to pinch Daisy and wrongs her by doing so. Even so, Daisy's killing him to prevent him from pinching her would be disproportionate. In *Fight or Flee*, Daisy can either escape before Tom arrives to kill her, or she can stand her ground and meet Tom's lethal threat with lethal defensive force. In this case, lethal defensive violence presumptively satisfies proportionality—a life for a life—but not necessity.

The proportionality condition is under particularly heavy scrutiny in what follows, though the other conditions enjoy significant roles as well. Escalatory cases, as well as cases involving conditional threats, raise sharp questions for our understanding of all three conditions on defensive violence, and can actually help to refine our understanding of them.

The main body of the article unfolds as follows. In Sect. 1, I outline the basic escalatory cases and fill in some contrasting positions. In Sects. 2 and 3, I examine the commitments of the proportionality condition. In Sect. 4, I consider the necessity condition in connection to these cases. In Sect. 5, I distinguish between two different types of escalation, and in Sects. 6 and 7 I examine conditional lethal threats. (The imminence condition will have an important bearing on what we should say

about these particular arguments). Section 8 states a couple of caveats and a brief conclusion.

1 Escalating Defence

I assume that Daisy's lethal defensive response to Tom in *Pinch* would be disproportionate, and therefore impermissible. A similarly structured case is *Steal*, in which Tom attempts to snatch Daisy's money.¹ Daisy's choices in *Steal* are either to repel Tom's attempted theft by killing him—she can frustrate it in no other way—or to acquiesce to the loss of her money. Killing him, however, would be disproportionate: although the money is Daisy's and Tom wrongs her by stealing it, she may not kill him to hold on to it.

Faced with only these two options in *Steal*, Daisy must simply say goodbye to her money. However, the story might unfold differently. Imagine, then, that Daisy does not face at the outset the simple choice between killing Tom and surrendering to the theft. She has the money, and can hold on to it—at least for the moment. In particular, Daisy might permissibly engage in sub-lethal defensive violence to protect her property. If rights against property are enforceable, then we should expect there to be such a level of permissible defensive violence that is proportionate to the harm imposed by Tom's initial action in *Steal*. Call that level of harm " H_{STEAL} ". For example, it might be permissible for Daisy to shove or punch Tom.² Of course, Tom might have set the initial aggressive threat at a much higher level, so that Daisy's defensive force of H_{STEAL} would be ineffective. These possibilities will be examined later, in Sect. 7. In the present case, Daisy's infliction of H_{STEAL} does enjoy some immediate efficacy against Tom's attempted wrongdoing.³

However we propose to fix the level of violence appropriate for H_{STEAL} , I assume that H_{STEAL} falls far short of H_{DEATH} , where H_{DEATH} denotes lethal harm. Defensive violence of H_{DEATH} would be justified only given a grave risk to Daisy's life or her central interests in bodily integrity.

Assuming that Tom has embarked on his attempted theft, we pick up the story where Daisy has inflicted defensive violence of H_{STEAL} on Tom and successfully retained her property. Tom might of course call the whole thing off at that point, which is what he ought to do. But suppose instead that he turns violent. He escalates the violence to H_{STEAL+} , which is a higher level of violence lying somewhere between H_{STEAL} and H_{DEATH} . In a renewed defensive response, Daisy now matches this level of violence with H_{STEAL+} . Faced with this enhanced level of defensive force, Tom now escalates all the way up to H_{DEATH} . To match Tom's aggressive threat of H_{DEATH} , Daisy now deploys violence of H_{DEATH} . Daisy and Tom have now locked themselves into a life and death struggle. If Daisy is successful in these defensive efforts, Tom will die. An attempted theft that was not worth the loss of

¹ There is a potential dissimilarity between them, however, which I will tackle in Sect. 3.

² The value of H_{STEAL} may also be sensitive to the value of the property targeted by the thief.

³ I will examine efficacy's connection to the necessity condition in greater depth in Sect. 4.

anyone's life at the start of this episode has now escalated into a violent struggle in which only one of the agents will be alive at the end of it. I call this variant of the original *Steal* case *Escalated Steal*.⁴

We can trace out the same sort of progression for a variant of *Pinch*, starting from a sub-lethal form of permissible defensive response of H_{PINCH} , but ending up in the same lethal place, through successive increments of $H_{\text{PINCH+}}$ and then H_{DEATH} . Call this analogous case *Escalated Pinch*. It does not matter for present purposes whether H_{STEAL} and H_{PINCH} state equivalent levels of harm. What matters is that, once Daisy has stepped on to this defensive escalator at an appropriately minor level of defensive violence, she may be heading towards the infliction of H_{DEATH} if Tom maintains his escalation of aggressive violence and she matches these escalations.

Is Daisy morally required to suffer a relatively small loss—the theft, or the pinch—in order to avoid this escalation to deadly violence? Or are her defensive escalations permissible, given Tom's moral responsibility for the escalating violence and the fact that, at any given stage, Daisy's violence is proportionate? After all, Daisy's defensive moves are either constituted by the sub-lethal level of harm (H_{STEAL} or H_{PINCH} , depending on the case) or by a series of increasingly violent defensive responses to Tom's escalating aggression against her. I think Daisy can permissibly escalate in both cases.⁵

One more preliminary point. Since there would be no escalation of violence but for Daisy's violent resistance, does that fact demonstrate in and by itself that Daisy is at least partly morally responsible for the harms that ensue? No, because this point over-generates implications. A causal partnership between defender and aggressor is instantiated by *all* defensive violence, which occurs due to the combination of aggressive violence and defensive violence. We do not routinely infer that *moral* responsibility for the harms generated by this violence must therefore be divided between the defender and the aggressor in such a way that the defender collects at least partial moral responsibility for it. If Daisy cannot permissibly escalate, it cannot be merely because it takes two to make a fight. It must be because, by doing so, she violates some more specific requirement. Which one, though?

2 Escalation and Proportionality

One of these more specific requirements is proportionality. Seth Lazar believes that *Escalated Pinch* “proves too much”, on the grounds that it supposedly demonstrates that “it is proportionate for the victim to kill the aggressor, ultimately, to prevent him from pinching her” (Lazar 2014: 30, 31). David Rodin concurs: “Defensive action is impermissible when it foreseeably produces harmful effects that are

⁴ Tom's aggressive escalation up to $H_{\text{STEAL+}}$ in response to Daisy's initial defensive violence of H_{STEAL} teaches us that what initially seemed to be a *theft* has evolved into a violent *robbery*. *Escalated Steal* is plainly a robbery, not a theft.

⁵ For arguments congenial to my conclusions, see Finlay (2015: 63–76) and McMahan (1994: 195–196). One of the aims in this paper is to flesh these suggestions out and provide replies to the critical concerns that arise.

disproportionate to the good one is seeking to achieve" (Rodin 2014: 82). The "good one is seeking to achieve", for Rodin, evidently consists in the prevention of the pinch. Is the Lazar-Rodin verdict a fair one? No, I think not. But we must proceed slowly to see why.

Lazar suggests that, in *Escalated Pinch*, a life is traded for a pinch: clearly disproportionate. This explains why, for him, it is impermissible for Daisy to escalate. But this claim about proportionality actually seems askew. Daisy deploys lethal violence to defend her life, not to prevent a pinch. At each stage, Daisy's defensive response is proportionate to the aggression Tom deploys against her. The same goes for *Escalated Steal*. In both these cases, Daisy matches Tom's aggressive responses at any given stage, which makes her defensive violence proportionate. Daisy's proportionality calculations are not obviously erroneous.

I suspect that Lazar misidentifies the ostensible source of the prohibition in *Escalated Pinch*. He writes: "Principles of self-defence constrain otherwise justified defence in ways that an unscrupulous attacker can manipulate" (Lazar 2014: 29). As far as it goes, Lazar's claim is true, and it explains why Daisy's resistance to Tom in the original version of *Pinch* must not be extravagantly violent. Unscrupulous aggressors can indeed price effective but proportionate defensive responses out of the market. Now Lazar thinks the same verdict carries over to *Escalated Pinch*. But for that verdict to stick, Lazar appears to be relying on the thought that, because Tom *will* escalate, the higher levels of defensive violence inflicted by Daisy later on must be disproportionate, even if at each level in the ensuing exchange they appear proportionate. So, because Tom *will* escalate further, Daisy lacks permissible defensive responses at these later times. Why should that be?

If Daisy deploys defensive violence equivalent to H_{PINCH} , then, by assumption, Tom will escalate in *Escalated Pinch*. Those who think it is impermissible for Daisy to defend herself against Tom must select a place for her to cease defensive violence. They seem to have two options. First, Daisy's violence can be no higher than H_{PINCH} , even if Tom then escalates to H_{PINCH+} . Second, Daisy is under the even more stringent requirement not to offer any violent resistance in the first place, which will rule out the possibility of justified escalation from the permissible starting point of H_{PINCH} . Both of these options seem unenticing. I consider them in turn.

The first option requires Daisy not to escalate beyond H_{PINCH} , regardless of what Tom does next. If Tom launched an attack on Daisy with the higher aggression of H_{PINCH+} , then Daisy could permissibly respond with H_{PINCH+} . But if Tom impermissibly escalates to H_{PINCH+} from an attack starting with H_{PINCH} , then Daisy lacks permission to match him at the higher level. In other words, Tom faces no permissible defensive resistance to aggression of H_{PINCH+} just as long as he deploys his aggression in escalated stages. If Tom can *build up* to H_{PINCH+} from a lower level of aggressive violence, then Daisy cannot permissibly defend herself with H_{PINCH+} . That implication is unsatisfactory.

Now for the second option, where Daisy is denied permission to step on to the defensive escalator in the first place. If Tom will escalate from H_{PINCH} to H_{PINCH+} , then Daisy cannot permissibly deploy a defensive response of even H_{PINCH} . A striking asymmetry presents itself here. If Tom does not in fact escalate, then Daisy's initial defensive response of H_{PINCH} counts as proportionate. There will be a

mini-struggle and the encounter, for better or worse, will be over. By contrast, if Tom does escalate, then Daisy's initial defensive response is disproportionate. It is only because Tom *will* escalate to a higher level of aggressive violence that Daisy's initial level of defensive violence is impermissible. Yet Tom's escalation is clearly *impermissible*. He never acquires justification for any violence at any point in the sequence. The permissibility of Daisy's defence is therefore being tied to the question of whether Tom will impermissibly escalate. This is odd. What is going on?

The commitments at work in this picture may reflect an interpretation of proportionality that does not tie proportionate response to one level at a time in the escalated violent sequence. To begin to see how this alternative interpretation of proportionality works, return to a segment of the escalated sequence in *Escalated Pinch*. As a reminder, Tom initially attempts to pinch Daisy. Then Daisy deploys H_{PINCH} against his attempt to pinch her. Faced with this obstacle, Tom next ascends to $H_{\text{PINCH}+}$, to which Daisy responds with $H_{\text{PINCH}+}$.⁶ Do Daisy's responses satisfy proportionality? At *each stage* in the overall sequence, proportionality *does* appear to be satisfied: Daisy's H_{PINCH} is proportionate to Tom's attempted pinch, and Daisy's $H_{\text{PINCH}+}$ is proportionate to Tom's $H_{\text{PINCH}+}$. Thus the ordinary notion of proportionality will make Daisy's responses permissible.

An alternative measure of proportionality seeks to compare levels of violence *across* these different times. In particular, Daisy's defence of $H_{\text{PINCH}+}$ is disproportionate to the original wrong threatened by Tom: namely, a pinch. We have already agreed that defensive responses to the original wrong posed to Daisy in *Pinch* can go no higher than H_{PINCH} . Now the violence has escalated to $H_{\text{PINCH}+}$ due to Daisy's refusal to submit to that initial wrong. By assumption, Daisy can see what is coming up. If she deploys H_{PINCH} at the earlier time, then Tom will escalate from H_{PINCH} to $H_{\text{PINCH}+}$, which will require, in turn, a defensive escalation to $H_{\text{PINCH}+}$. The charge lying in wait for Daisy is that $H_{\text{PINCH}+}$ is a disproportionately excessive level of defensive violence to apply to the offence that originally triggered the conflict between Daisy and Tom.

Call this alternative measure of proportionality "intertemporal proportionality". Intertemporal proportionality measures proportionality over a series of chronologically adjacent discreet acts of violence between two agents, such that the proportionality of the defender's responses is measured by the severity of the aggressor's initial act of aggression. In constructing its verdicts on the proportionality of defensive response, the test for intertemporal proportionality mixes the times in the escalating sequence. As applied to the itinerary of *Escalated Pinch*, the yardstick for the proportionality of Daisy's defensive responses is fixed by Tom's original aggression of H_{PINCH} . Because Daisy's later defensive response of $H_{\text{PINCH}+}$ exceeds H_{PINCH} , it is disproportionate.

⁶ Two further notes: first, I am assuming that the sequence in *Escalated Pinch* unfolds in the same way as *Escalated Steal*; second, I need to attend to only a fragment of the overall escalated sequence in order to explore the underlying ideas.

Intertemporal proportionality disregards other parts of the original comparison between Tom's actions and Daisy's actions, and this is problematic. Two parts of it, in effect, are left out.

First, intertemporal proportionality effectively ignores Tom's escalation to H_{PINCH+} . In *Escalated Pinch*, Daisy responds with H_{PINCH+} to Tom's escalation to H_{PINCH+} . By assumption, if Tom did not ascend to H_{PINCH+} , Daisy would not respond with H_{PINCH+} .⁷ If we can condemn Daisy for inflicting H_{PINCH+} , because H_{PINCH+} is disproportionate relative to Tom's original aggression of H_{PINCH} , then to all intents and purposes the intertemporal proportionality test overlooks Tom's escalation from H_{PINCH} to H_{PINCH+} . If we had taken Tom's escalation to H_{PINCH+} into consideration, then Daisy's defensive escalation to H_{PINCH+} would be permissible, because it would satisfy proportionality at that escalated level.

Second, intertemporal proportionality ignores the context for Daisy's original defensive response of H_{PINCH} . On its preferred method of moral book-keeping, Daisy deploys H_{PINCH+} against Tom's H_{PINCH} . Intertemporal proportionality identifies the disproportionateness of Daisy's response by pointing to a mismatch between Tom's H_{PINCH} and Daisy's H_{PINCH+} . This verdict ignores both the fact that Daisy's original defensive response of H_{PINCH} did satisfy proportionality, and the fact that Daisy does not escalate from H_{PINCH} to H_{PINCH+} in response merely to Tom's original threat of H_{PINCH} . Daisy's H_{PINCH+} is plainly a response to Tom's escalation from H_{PINCH} to H_{PINCH+} . The complaint that Daisy disproportionately applies H_{PINCH+} to aggression that permits only H_{PINCH} overlooks the fact that Tom impermissibly escalates from H_{PINCH} to H_{PINCH+} .

So why ignore Tom's escalated aggression and Daisy's initial defensive response? Why pretend, in effect, that Daisy's defensive response to the original threat embarks from the higher level of H_{PINCH+} ? Since it does not have obvious answers to these questions, intertemporal proportionality is a dubious resource.⁸

3 Property Versus Lives

In *Escalated Steal*, Tom escalates his violence, and Daisy matches it in defensive violence. At some point, what originally started out as resistance to a theft escalates, so that Daisy must eventually kill Tom in order to defend herself. Daisy's reasons for escalating are complex, and fall into two categories. First, there are *purely defensive* reasons, concerning Daisy's responses to Tom's higher levels of aggression against her and her bodily integrity. Second, there are *property-focused* reasons, concerning her refusal to surrender her money. The risk here is that the purely defensive reasons interact poorly with the property-focused reasons, and this introduces a potential disanalogy between *Escalated Steal* and *Escalated Pinch*. The aim of this section is to investigate this potential disanalogy.

⁷ I will consider other possibilities in Sect. 5.

⁸ We shall see, in Sect. 4, whether an intertemporal necessity condition fares any better. Short answer/spoiler: no.

In what I will call the *Continued Possession Variant* of *Escalated Steal*, Daisy still has her money, and can evade Tom's continued aggression against her by handing it over to him. But she chooses to fight instead, despite the foreseeable escalation of violence. I maintain that Daisy's infliction of H_{DEATH} in the Continued Possession Variant is permissible. In the *Reclaiming Variant* of *Escalated Steal*, Daisy has her money stolen, and then renews her fighting with Tom, inflicting H_{DEATH} on him to reclaim it. Daisy's escalation in the Reclaiming Variant is clearly impermissible because this case is essentially no different from the original *Steal*. But if Daisy can permissibly inflict H_{DEATH} in order to *hang on to* her money, as in the Continued Possession Variant, then why can't she permissibly inflict H_{DEATH} in order to *regain* it, as in the Reclaiming Variant? These two variants may appear to stand or fall together.

If there are no deep moral differences between the Continued Possession Variant and the Reclaiming Variant, then permitting Daisy to kill Tom in the Continued Possession Variant will also permit her to kill Tom in the Reclaiming Variant. The rest of this argument now works by *modus tollens*. Daisy's violent response in the Reclaiming Variant is impermissible: it would be disproportionate for Daisy to kill Tom to thwart his theft. But if the Reclaiming Variant is morally objectionable, the Continued Possession Variant must also be morally objectionable. I call this the "Analogy Argument". It rests upon an analogy between the Reclaiming Variant and Continued Possession Variant, and it uses the impermissibility of the former to establish the impermissibility of the latter.

As I see it, the Analogy Argument fails. The Reclaiming Variant and the Continued Possession Variant are relevantly different. Even though both variants reflect Daisy's broad interest in hanging on to her money, they are morally distinguishable. In particular, the Analogy Argument ignores the fact that, while she still the money, Daisy is threatened by Tom's actions. Tom inflicts harm on Daisy in order to take her money. His ultimate end may be only theft, rather than theft-and-assault, but the means he employs to that end involve personal harm to Daisy. (In other words, Tom's ultimate aim may be theft, but he achieves it by violent robbery.) This aspect of the wrong makes the all-important difference. Its significance is that Daisy's property-focused reasons cannot fail to trigger her defensive reasons as well.

In the Reclaiming Variant, Tom seizes Daisy's money. He ought not to have done that, but Daisy can only permissibly inflict H_{STEAL} on him to reclaim her money. Now assume in the Continued Possession Variant that we have reached the stage where Tom is threatening Daisy with $H_{\text{STEAL}+}$ if she does not give him her money. The crucial question is whether, in the Continued Possession Variant, Daisy can permissibly respond with $H_{\text{STEAL}+}$.

We cannot reasonably deny permission to Daisy to deploy the more violent $H_{\text{STEAL}+}$. As things stand, given the fact that Daisy is non-labile, Tom has no right to harm Daisy for *any* immediate reason,⁹ least of all to rob her. Tom's robbery of Daisy only increases the ways in which he is wronging her. Tom cannot permissibly

⁹ I leave aside cases involving the lesser evil justification or enforcement of duties of rescue. In *Escalated Steal*, Tom is engaged in robbery, and such considerations do not apply to him.

take Daisy's money, *and* he cannot permissibly harm Daisy. *Escalated Steal* combines elements of both wrongs. In an ordinary case of self-defence in which Tom's threat to her is equivalent to $H_{\text{STEAL}+}$, Daisy may defend herself with proportionate force of $H_{\text{STEAL}+}$. In a purer case of theft in which only Daisy's property is threatened, she may deploy force of only H_{STEAL} . The Continued Possession Variant combines both the wrongs of pure theft and the wrongs of danger to bodily integrity. Daisy's property-based reasons and her defensive reasons are thus both at issue, because both her bodily integrity and her money are under threat. So why is Daisy permitted to deploy only H_{STEAL} and not $H_{\text{STEAL}+}$? Daisy is non-labile; she would be wronged by Tom's attack on her. The reason why Tom attacks her is to steal her money. The fact that Tom undertakes this unjust attack on Daisy in order to commit a *further* wrong—stealing from her—cannot improve his moral position, and it cannot worsen Daisy's defensive position. Thus Tom's violent attempted robbery of Daisy does not reduce Daisy's defensive rights.

Daisy's proportionate defence permits a level of harm of $H_{\text{STEAL}+}$. But $H_{\text{STEAL}+}$ exceeds H_{STEAL} . So there is a morally relevant difference between the Continued Possession Variant and the Reclaiming Variant that justifies different levels of permissible defensive violence. This confirms, in turn, that *Escalated Steal* is analogous to *Escalated Pinch*.

4 Necessity, Futility, and the Duty to Retreat

On the face of it, the burdens of the necessity condition in its standard interpretation fall entirely on Tom, not Daisy. Each of them *can* retreat, but the necessity condition protects only Daisy, not Tom. The burdens and advantages induced by the necessity condition, on this way of understanding it, are sharply asymmetrical. At each level of response, Daisy is entitled to protect herself at least cost, and she thus satisfies necessity, assuming that she cannot avoid this particular threat. The retreat requirement is only a duty to retreat from avoidable unjustified violence. A defender is under no duty to acquiesce to unjustified aggressive violence in order to avoid exposure to even higher levels of unjustified aggressive violence that would require, in turn, higher levels of defensive violence. At each level of the escalation, necessity will not stop being a reliable ally of Daisy. By contrast, at each level, Tom lacks the right to harm Daisy, and is under a standing instruction to desist. Tom ought to retreat, but no such instruction applies to Daisy. If Daisy enjoys the choice between retreat and defence, she should retreat. However, if retreat is not possible, necessity permits Daisy to defend herself at least cost.

An immediate caveat: more sophisticated conceptions of necessity have something to say about risk and the probabilities of success, as well as the proper division of harm among defenders, aggressors, and bystanders. I will not engage with issues of probability here, and the two-agent structure of these cases permits me to ignore the implications for bystanders. (For more on these issues, see, for example, Lazar 2012; McMahan 2015; Oberman 2020; and Quong 2020, ch. 5.) But even within the limitations of my two-agent framework, a couple of further challenges should be discussed. The first of these challenges concerns futility, and the second of them

suggests a re-division of burdens between aggressor and defender that draws upon an intertemporal version of necessity. I discuss them in turn.

First, futility. In *Escalated Pinch*, Tom will escalate if Daisy resists his initial aggression. The fact that Tom will escalate raises in turn the possibility that Daisy's initial level of defensive response, H_{PINCH} , will be *futile*. It will not stop Tom in his tracks, but merely prompt him to escalate further. Even though there is no disproportionateness as such between Tom's aggression of H_{PINCH} and Daisy's defence of H_{PINCH} , Daisy's defensive violence lacks efficacy. It merely serves as a basis for Tom's later escalation to $H_{\text{PINCH}+}$. Once efficacy is lost, then Daisy's permission to engage in it is also moot. The necessity condition may wish to uphold the permissibility of defensive violence as a last resort, but, as Daniel Statman reminds us, "[f] or some course of action to be a last resort, it must first be a resort" (Statman 2008: 663).

There is more than one way of tackling this worry about futility. First, we may think that even non-efficacious defensive violence is permissible if it contributes to the defender's self-respect or dignity (Statman 2008). I have some sympathy with this line of argument.¹⁰ But there is no need to rely on it for *Escalated Pinch*, because in my view the charge of futility is misplaced. My second point seeks to vindicate this claim. Because the violence escalates on both sides, there is a robust sense in which Daisy's initial defensive response of H_{PINCH} is efficacious. After all, Daisy has ensured that Tom does not wrong her by *pinching* her. The fight is far from over, as we know, but Daisy has in fact successfully defended herself from the initial threat. She still risks Tom's harming her at higher levels of violence, but she no longer has to worry about the charge of futility. Whatever happens next, Daisy will have stopped Tom from pinching her, which was her original defensive concern.¹¹

We can grasp this point more vividly by considering another series of encounters:

Weekend Escalation: On Friday, Tom tries to pinch Daisy. If Daisy successfully resists Tom with proportionate force, then, on Saturday, Tom will try to break Daisy's arm. If Daisy successfully resists Tom's Saturday attack with proportionate force, then, on Sunday, Tom will try to kill Daisy. To defend herself, Daisy will have no choice but to deploy lethal force against him.

In *Weekend Escalation*, the violence unfolds in a series of separate, temporally non-adjacent conflicts. Yet the deep conditional structure in *Weekend Escalation* is shared with *Escalated Pinch*. Just as, in *Weekend Escalation*, Daisy's reward for resistance on Friday is further violence on Saturday, Daisy's earlier resistance in *Escalated Pinch* is met by greater violence later on. What *Weekend Escalation* makes plain, though, is that Daisy's defensive response to the Friday attack

¹⁰ Though see Frowe (2015) for a good critical discussion.

¹¹ A similar point about the efficacy of defensive action in cases where a defender is confronted by multiple aggressors is made by Frowe (2014: 111): even if a defender's defensive action against, say, two aggressors fails to protect her against assault from the *other* aggressors, it will protect her from *these* aggressors, which is enough to justify defensive violence against *them*.

cannot be reasonably condemned as inefficacious, on the grounds that Tom will then attempt to break her arm on Saturday. She remains unpinched on Friday, even if she now needs to start worrying about her arm. And similarly, Daisy's defensive response to the Saturday attack, where she defends herself against Tom's attempt to break her arm, cannot be condemned as inefficacious on the grounds that Tom will attempt to kill her on Sunday. Daisy's arm is safe. It is her life that is now in danger.

Two further comments. First, if Daisy later suffers a worse injury, and this is in part due to her resistance to Tom's earlier aggression, she is unlikely to console herself with the thought that, after all, these earlier defensive efforts avoided *futility*. But it does not follow that these earlier efforts were *impermissible*, even if, by the end of the weekend, they turn out to have been less than sensible because they only got her into worse sorts of trouble. At the time, they were permissible acts of defence against specific unjustified threats. We need to distinguish between the *moral permissibility* and the *prudential wisdom* of Daisy's actions. These issues should not be conflated. Second, the suspicion might arise that Daisy's defensive permissions in *Weekend Escalation* are more secure than they are in *Escalated Pinch* simply because the gaps between Tom's attacks are longer. In *Weekend Escalation*, the attacks are all one day apart. When the temporal distances shrink, as they do in *Escalated Pinch*, the basis for Daisy's defensive resistance may seem less secure. In my view, we should resist this suspicion. Imagine a variant of *Weekend Escalation*, which I will call *Friday Night Escalation*. Here Tom's attacks are less spaciouly timetabled to fall on Friday night only: he will attempt to pinch her at 6 p.m., to break her arm at 8 p.m., and to kill her at 10 p.m. As I see it, Daisy's defensive options are not weakened in *Friday Night Escalation* because Tom's projected attacks are on a tighter schedule. If she can perform the same actions, with the same defensive results, as she did in *Weekend Escalation*, she may do so, and for the same basic reasons.

As a result of her defensive resistance, Daisy is unharmed after the first two encounters in *Weekend Escalation*. She is now ready to fight—and perhaps to die—another day. Nonetheless, the charge of futility is misplaced. The same lesson applies to *Escalated Pinch*.

The second challenge concerns intertemporal necessity. I start with the following case, due to Jeff McMahan:

Kill or Allow Bruising: Daisy will be killed by Tom unless she either (1) kills Tom, leaving her completely unharmed, or (2) incapacitates him without harming him, in which his attempt to kill her will leave her with a painful bruise (McMahan 2015: 189).¹²

The notion of necessity we have operated with so far permits Daisy to select (1) rather than (2). One might think that a more sophisticated version of necessity will prescribe (2) rather than (1), and that this version of necessity is preferable because it offers tighter management over permissible defensive violence, leading to less morally weighted harm overall between the liable aggressor and the non-liable defender. If this is so, then the challenge to *Escalated Steal* and *Escalated Pinch* will be that

¹² I take the name of the case from McMahan but have altered the names of the agents.

Daisy is required to accept the smaller harm (the theft, or the pinch), in order to save Tom from the greater harm of death when the escalations have run their course.

On the face of it, the analogy between *Kill or Allow Bruising* and *Escalated Pinch* is not secure.¹³ In *Kill or Allow Bruising*, there is one threat—Tom's deadly threat against Daisy—that needs to be dealt with. Adequate defence against this threat requires either Tom to die or Daisy to accept a painful bruise. In *Escalated Pinch*, there is also one threat at the start: the pinch. Daisy deals with this threat by deploying minor defensive violence: H_{PINCH} . By assumption, the harm to Daisy if Tom's pinch goes unchallenged is no less than H_{PINCH} . Choosing not to contest Tom's pinch is therefore not analogous to the choice of (2) over (1) in *Kill or Allow Bruising*. Moreover, since the escalated sequence of aggressive harms allows each further aggressive harm to be dealt with as it comes, none of Daisy's proportionate responses can be immediately likened to the choice between (1) and (2) in *Kill or Allow Bruising*, even at the higher levels of violence. When Tom ascends to $H_{\text{PINCH}+}$, Daisy responds with $H_{\text{PINCH}+}$, and when Tom escalates to H_{DEATH} , Daisy responds with H_{DEATH} . These are all cases of parity. No savings in harm are in the offing. As long as we examine each level of escalated violence separately, necessity is satisfied. Once again, necessity will be a reliable ally of Daisy rather than Tom just as long as we respect the escalated layers of violence.

The only way of bringing *Escalated Pinch* into the conceptual orbit of *Kill or Allow Bruising* is to invest in an intertemporal conception of necessity, which keeps a more watchful eye on the transition between the levels of escalated violence. Intertemporal necessity is analogous to intertemporal proportionality. The standard for intertemporal necessity is set by the aggressor's earlier level of violence in an escalated sequence of discreet violent encounters between the aggressor and the defender. Since the defender can avoid exposure to higher levels of aggression by declining to resist the aggressor's lowest level of aggression, the defender's escalation of violent defensive responses risks failing the interpersonal necessity test.

Let us see how intertemporal necessity applies to *Escalated Pinch*. Tom's aggression ascends from H_{PINCH} to $H_{\text{PINCH}+}$ if Daisy responds to Tom's attempted pinch with H_{PINCH} . It follows that Daisy could have avoided the risk of exposure to $H_{\text{PINCH}+}$ if she had refrained from countering Tom's pinch with H_{PINCH} . In this sense, then, Daisy's exposure to Tom's $H_{\text{PINCH}+}$ is avoidable. She could have avoided it by exercising a retreat option at an earlier point by declining to resist Tom's pinch with H_{PINCH} . Daisy has a retreat option, then, but it can be exercised only pre-emptively.

The claim that intertemporal necessity should fix Daisy's defensive permissions may be able to draw further support from McMahan's *Kill or Allow Bruising*. This is because *Kill or Allow Bruising* can be regarded as a temporally compressed form of *Escalated Pinch*. The difference between them is that, in *Kill or Allow Bruising*, Tom delivers his threat all at once, and Daisy is required by McMahan to absorb a minor harm (a bruise) in order to spare Tom from a much greater harm (death). This instruction can perhaps be likened to the demand, in

¹³ I cannot explore here whether McMahan's case carries force for other necessity judgments. My purpose here is only to establish that my commitments are unchallenged by *Kill or Avoid Bruising*.

Escalated Pinch, that Daisy refrain from resisting the lesser harm (the pinch) in order to spare Tom from suffering the greater harm later in the sequence (death). In this sense, the requirement that Daisy not escalate is governed by the reasoning McMahan applies to *Kill or Allow Bruising*. *Escalated Pinch* simply extends McMahan's reasoning in an intertemporal rather than intratemporal way.

In my view, the analogy between *Kill or Allow Bruising* and *Escalated Pinch* is less secure than it may first seem. The relevant difference between them is that Tom's threat in *Kill or Allow Bruising* is basically a *fait accompli*: Tom has already issued this threat, and it is now Daisy's task to manage it. At *this* point, there is no further decision Tom has to take. Dividing the harms between him and Daisy is now instead an unavoidable distributive problem *for her*. McMahan's claim is that Daisy successfully manages this problem only by minimizing the total amount of defensive weighted harm. Even though their interests are weighted differently because Tom is liable and Daisy is non-labile, a combination of <a bruise to Daisy, no harm to Tom> is morally preferable to the combination of <no harm to Daisy, death to Tom>. The escalated violence in *Escalated Pinch*, by contrast, is *not a fait accompli*. Tom is not *fated* to escalate. If Daisy is required in *Escalated Pinch* to submit to the pinch in order to deter Tom from escalation, then her task is not just to manage harm that Tom has already issued, but to appease or manage *Tom* by sacrificing herself. This goes beyond any lessons that we can recover from *Kill or Allow Bruising*. *Kill or Allow Bruising* does not tell Daisy to absorb some degree of harm in order to get Tom to refrain from renewing his aggression in strengthened form.

Here is another way of expressing the point. Intertemporal necessity insists that Daisy wholly absorbs a minor aggressive harm in order to spare Tom a more serious defensive harm that will be inflicted only if Tom decides, impermissibly, to escalate. Tom's responsibility for that escalation is ignored. In *Kill or Allow Bruising*, by contrast, Tom's violence is not ignored. It is too late to ignore it. Daisy's task is simply to manage it by minimizing morally weighted harm. It is entirely consistent with Daisy's permission to escalate in *Escalated Pinch* that she should be governed by this further defensive scruple from McMahan.

Once this disanalogy between the two cases is noted, we are free to criticize intertemporal necessity in ways that parallel the criticisms of intertemporal proportionality. Like intertemporal proportionality, intertemporal necessity annuls defensive options earlier in the sequence, which counts against it. Because Tom *will* escalate to aggressive violence of $H_{\text{PINCH}+}$, intertemporal necessity instructs Daisy not to respond with H_{PINCH} . It follows that intertemporal necessity cancels Daisy's initial defensive options, even though these are perfectly proportionate to Tom's various acts of aggression.

Why should a revamped necessity condition collude in this way to ignore what Tom has done, and will do, when what he has done and will do are impermissible? These are serious concerns about intertemporal necessity. In my view, Daisy does not face decisive discouragement from the necessity condition.

5 Jumping Ahead: Imminence and Conditional Threats

In both *Escalated Pinch* and *Escalated Steal*, it is Tom who initially escalates at each level of escalation, and Daisy who responds with an enhanced defensive response. These cases involve what Lazar calls “coercive” conditional threats.

In coercive threats, it is the aggressor, not the defender, who sets the pattern of escalation. If Daisy resists, Tom will escalate, but Daisy will not unilaterally escalate. It is Tom who leads the escalation, while Daisy responds to it.

Coercive conditional threats contrast with “defensive” conditional threats. In defensive conditional threats, it is the defender, not the aggressor, who sets the pattern of escalation. If Daisy escalates, Tom will escalate in response, but Tom will not unilaterally escalate (Lazar 2014: 26). It is Daisy who drives the escalation to higher levels of violence, even if we can count on Tom to match these defensive escalations with aggressive escalations.

Both types of threat are *conditional* in the sense that, if Daisy provides any resistance, then Tom will either escalate his aggressive threat (in coercive cases), or match Daisy’s defensive escalation with an equivalent escalation of aggressive force (in defensive cases). Correlatively, if Daisy does not resist (in coercive cases) or unilaterally escalate (in defensive cases), then she will not face escalated aggression from Tom. Conditional threats do not merely specify threats that might be coming the defender’s way, but also indicate to the defender certain means of avoiding them. Whether conditional threats transpire into actual threats is, in part, up to the defender. That is what makes them, in the relevant sense, conditional.¹⁴

I have argued that Daisy can justifiably match Tom’s escalations with defensive escalations in cases involving coercive conditional threats, such as *Escalated Pinch* and *Escalated Steal*. The next question is this: can Daisy take the lead instead of Tom? Can Daisy’s defensive resistance permissibly take the form of a *defensive* conditional threat? Daisy might believe that, if Tom escalated first, then she could permissibly escalate. But she may then be seized by the thought that there is no need for her to wait for him to do that. Defensive violence need not be purely reactive. The imminence condition does not require defenders to wait until they are actually under attack before defending themselves. It may be too late by then; any conceivable defensive advantage may have been lost. Permissible defensive violence can thus pre-empt aggressive violence, not just follow in its wake, grimly hoping for the best.

These reflections on the imminence condition appear to establish that, when lethal defensive force is applied, the defender’s violence may permissibly exceed the aggressor’s violence. Daisy may thus suppose that she can permissibly escalate via a defensive conditional threat. Call this variant of the case *Pre-emptive Escalated Steal*. (The analogous case for *Escalated Pinch* is *Pre-emptive Escalated Pinch*.)

¹⁴ Of course, other circumstances might play a role in fixing whether the aggressor actually delivers on his threat. His threat may be dependent, for example, on whether it is raining on the day in question. But this sort of contingency does not matter to conditional threats: in the relevant sense, a conditional threat is a threat whose activation is crucially dependent on what the defender chooses to do in response to it.

As we know, Lazar judges it impermissible for Daisy to escalate even when confronted by coercive conditional threats: Daisy may therefore not escalate in *Escalated Steal*. Unsurprisingly, he takes the same attitude to defensive conditional threats, so that this prohibition extends to *Pre-emptive Escalated Steal*. But he also contends that defensive conditional threats are marred by a further problem, to which we need to pay greater attention.

In particular, Lazar argues that this “argument from preventive defence” is circular. Imagine that Daisy is wondering at T_0 whether she will be justified in applying lethal defensive force to Tom at T_1 , given the fact that she will thereby be pre-empting his lethal aggressive force at T_2 . Lazar writes:

[T]he argument from preventive defence justifies lethal defence at T_1 by assuming (at T_0) that the threat posed by the aggressor at T_2 will be unjustified. To show that the T_2 threat is unjustified, we must show that the T_1 threat is justified; but to show that the T_1 threat is justified, we must assume that the T_2 threat is unjustified (Lazar 2014: 28).

Lazar argues that two conditions need to be satisfied in order for Daisy’s pre-emptive lethal defence to be justified, and that they cannot be satisfied simultaneously. First, the permissibility of Daisy’s violence at T_1 depends upon the impermissibility of Tom’s violence at T_2 (this is the “aggression condition”). Second, the impermissibility of Tom’s violence at T_2 depends upon the permissibility of Daisy’s violence at T_1 (this is the “defence condition”). The defence condition cannot be satisfied unless the aggression condition is satisfied, and the aggression condition cannot be satisfied unless the defence condition is satisfied. When we combine these two conditions, the justificatory possibilities evaporate. Neither condition can obtain a foothold in the absence of the other condition’s obtaining a foothold.

Now we already know that Tom’s escalation at T_2 would be unjustified, since this is just one sorry episode in a sequence in aggressive violence that is impermissible at each step, and impermissible overall. But this must mean that the aggression condition is satisfied and that the justificatory circle is now broken: mustn’t it? Lazar anticipates this likely reply. His response reveals that the defence condition is a much tougher condition to meet. It is the defence condition that is really calling the shots in this argument. In the following passage from him, I substitute “Daisy” and “Tom” for “A” and “B”, respectively:

If a threat posed by Tom at T_2 will eventuate only if Daisy acts unjustifiably at T_1 , *then even if that future threat will be unjustified on independent grounds*, it remains inadmissible in Daisy’s proportionality calculation at T_0 because it is conditioned on [her] intervening wrongful action at T_1 . If that action is wrongful, [she] ought not to do it, and the T_2 threat will not eventuate (Lazar 2014: 28; emphasis added).

We need to distinguish between different possibilities. If Tom will definitely, come what may, deploy lethal aggressive force at T_2 , then Daisy can permissibly pre-empt that threat at T_1 with proportionate defensive force. The imminence

condition has no quarrel with this type of pre-emptive defensive force. In *Pre-emptive Escalated Steal*, however, Tom's later lethal violence is conditional on Daisy's escalated defensive violence. By assumption, without prior defensive escalation, Tom will not escalate.

This factor may make a moral difference. It suggests that Daisy still needs an independent justification for escalating. If Daisy did not escalate, then Tom would not escalate. It follows that Daisy lacks a specific justification for pre-empting Tom's aggression by escalating first. I am inclined to think that Lazar is right about this case. But another type of case has a more unsettled status. I shall describe it in the next section and argue that the reasoning Lazar applies to *Pre-emptive Escalated Steal* does not apply to it.

6 Modestly Pre-emptive Cases

To recapitulate: in *Pre-emptive Escalated Steal*, Daisy leads the escalation, while Tom responds, and in *Escalated Steal*, Tom leads the escalation, while Daisy responds. A third case combines aspects of both cases. In *Modestly Pre-emptive Escalated Steal*, Daisy knows that Tom has a prior commitment to coercive aggressive escalation. Tom conforms to the profile of the aggressor in a coercive conditional case, not a defensive conditional case. Tom's disposition in *Modestly Pre-emptive Escalated Steal* is thus identical with his disposition in *Escalated Steal*, but not *Pre-emptive Escalated Steal*.

Daisy's question is now this: rather than simply matching Tom's aggressive escalations and waiting for the escalations to play out so that the lethal climax is reached in due course, can she permissibly omit the intermediate steps and ascend immediately to lethal defensive violence? Can she go to H_{DEATH} immediately, rather than ascending to H_{DEATH} , if she has to, in incremental steps?

There are different axes on which we can evaluate the similarity or dissimilarity among these cases. In the first instance, we might focus on either *dispositions* or *actions*.

As far as Tom's disposition is concerned, *Modestly Pre-emptive Escalated Steal* is more similar to *Escalated Steal* than to *Pre-emptive Escalated Steal*. Tom is disposed to engage in coercive conditional threats in both *Escalated Steal* and *Modestly Pre-emptive Escalated Steal*. Under this dimension of assessment, it is *Pre-emptive Escalated Steal* that is the outlier, because Tom's disposition in this case is to engage in only a defensive conditional threat.

Now for actions: *Modestly Pre-emptive Escalated Steal* seems nearer to *Pre-emptive Escalated Steal* than to *Escalated Steal*. In both *Modestly Pre-emptive Escalated Steal* and *Pre-emptive Escalated Steal*, Daisy escalates to lethal levels of violence ahead of Tom. In *Escalated Steal*, by contrast, Tom reaches a lethal level of aggressive threat ahead of Daisy.

I have already argued for the permissibility of Daisy's escalation in *Escalated Steal*, and for its impermissibility in *Pre-emptive Escalated Steal*. What verdict should we reach for *Modestly Pre-emptive Escalated Steal*? Should we centre critical

operations on dispositions or actions, or reject the relevance of this dichotomy for arriving at the correct answer?

In my view, Daisy can permissibly escalate in *Modestly Pre-emptive Escalated Steal*. This is not due squarely to the fact that Tom's disposition in *Modestly Pre-emptive Escalated Steal* is shared with his disposition in *Escalated Steal*. Sameness of disposition has a role to play, but it is embedded in a more complex explanation. There are in fact four relevant considerations, which combine to produce a permissibility verdict. I shall now enumerate them.

First, and as before, the imminence condition is in play. The mere fact that, for some arbitrarily chosen level of violence, H^* , Daisy anticipates Tom's violence by deploying defensive violence equivalent to H^* before he unleashes counter-defensive violence equivalent to H^* does not render her defensive response impermissible. Now this implication of the imminence condition is not enough to make it permissible for Daisy to escalate in *Modestly Pre-emptive Escalated Steal*. It was not enough, after all, to save *Pre-emptive Escalated Steal*. But it can still play a role in the overall justification when partnered with other relevant considerations.

Second, and due to Tom's disposition in *Modestly Pre-emptive Escalated Steal*, Daisy knows that, if she embarks from a defensive response of H_{STEAL} , Tom will escalate to H_{STEAL+} and then to H_{DEATH} . In this case, it makes no difference whether Daisy deploys the less violent H_{STEAL} , or the more violent H_{DEATH} . Either way, Tom will eventually counter-defend himself with H_{DEATH} , if Daisy defends herself with any degree of proportionate force.¹⁵

Third, it would at least be permissible—or so I have argued—for Daisy to deploy H_{STEAL} in all these cases, including *Modestly Pre-emptive Escalated Steal*. But a defensive response of H_{STEAL} elicits the same deadly counter-defensive response as the more violent defensive response of H_{DEATH} , at least when the sequence is played out in full.¹⁶ This suggests that Daisy can permissibly perform the more violent defensive act now; she can escalate to H_{DEATH} immediately. Yes, she reaches this level of violence ahead of Tom, but the imminence condition does not condemn this modest form of pre-emptive defensive violence if Tom was going to get there anyway through unilateral aggressive escalation. Tom's escalation to lethal levels is a foreseeable result of Daisy's stepping on to the defensive escalator at any stage—which, again, Daisy was permitted to do.

Fourth, there is no sense in which Daisy's enhanced pre-emptive defence constitutes a form of *entrapment* of Tom, where he is placed in circumstances inducing his performance of a morally wrongful act that could and should have been avoided. By assumption, Tom's disposition is such that he will escalate to lethal levels of aggression regardless of whether Daisy deploys the lesser or the greater defensive violence. She was definitely entitled to deploy lesser violence, and that would have led Tom to a lethal encounter in due course.

¹⁵ Note that this "either way" condition does *not* hold in *Pre-emptive Escalated Steal*, since in that case Tom will reach lethal levels of aggression *only if* Daisy escalates first.

¹⁶ As before, the sequence in full will go from H_{STEAL} to H_{STEAL+} to H_{DEATH} .

It follows from these points, as I see it, that Daisy is permitted to deploy the greater violence, since the only real difference here is that the lethal engagement between them is reached sooner.

A final note and major caveat before moving on. I have relied, in *Modestly Pre-emptive Escalated Steal*, on a highly demanding epistemic assumption: Daisy knows that Tom has the disposition associated with coercive conditional threats and that he will not lose it in the course of their confrontation. This is *not* knowledge that Daisy would ordinarily have. Even if Tom's moral track record was poor, he might change his mind. He might switch, without prior announcement, to posing only a defensive conditional threat. (Bad enough, but not as bad as posing a coercive conditional threat.) For most ordinary contexts, then, Daisy should treat this attack *as if* Tom is posing a defensive conditional threat, not a coercive conditional threat, and she should exercise due caution by not unilaterally pre-empting the escalation. *If* the (demanding) epistemic conditions are met, however, then Daisy acts permissibly. This is a significant result, even given these demanding epistemic assumptions, and I will draw on it in the next section.

7 Lethal Conditional Threats

One prominent reason for Lazar's dismissiveness towards *Escalated Pinch* is that escalated cases are irrelevant to the analysis of most ordinary violent encounters (Lazar 2014: 31). Ordinary violent encounters often contain conditional threats, and most cases of conditional threat are structured in ways that dispense with escalated sequences. These cases conform to the structure of the familiar threat issued by the stock highwayman's "Your money or your life!" (*Highwayman*). The conditional threat in *Highwayman* presents the defender, in effect, with a three-way choice: she can part with her money; or she can refuse to part with it and die; or she can fight back with perhaps lethal violence and subdue the robber. But this three-way choice does not yield escalatory possibilities as such.

Consider another case of conditional deadly threat, similar in structure to *Highwayman*, but featuring our usual agents. In *Lethal Conditional Steal*, Tom conditionally threatens to kill Daisy unless she hands over her money. There is no escalation of violence in this case. No sequence takes Daisy from permissible minor violence to permissible major violence. Daisy cannot earn the right to apply larger levels of defensive violence against Tom by initially entering the violent exchange between them at a lower and more easily justifiable level of violence, and then ascending in steps of greater violence, each of which—or so I have argued—is justified. She appears to lack initial access to the conditions of that escalation-involving justification.

There is a way forward. In fact, I will consider two possibilities. In what I call the *Augmentation Argument*, I apply the deep structure of *Modestly Pre-emptive Escalated Steal* to *Lethal Conditional Steal* in order to deliver Daisy's defensive permissions. But before I do that, I will more briefly consider the *Culpability Argument*.

According to the Culpability Argument, Tom is not merely attempting to steal from Daisy in *Lethal Conditional Steal*. He demonstrates a willingness to kill her

unless she complies with his demand to surrender her money. That makes him morally worse than an ordinary thief. He is demonstrating that he is, in effect, a violent robber, not just a thief. This lethal conditional threat considerably increases his culpability. Thomas Hurka and Gerhard Øverland have independently argued that this robber's heightened culpability may augment the defender's defensive rights against him. Indeed, the greater culpability of this conditionally deadly aggressor might enhance his liability to such a degree that the defender can permissibly kill him (Hurka 2005: 54–55; Øverland 2010: 334–340).

The Culpability Argument need not be cast in competition with the Augmentation Argument, which I explore next. On the face of it, the Culpability Argument may offer merely an alternative route to the same conclusions. Yet it seems to me that the justificatory route set out by the Augmentation Argument is the deeper one. First, it offers more detail on the mechanics of the relevant normative claims. Second, it may be of greater interest to those who deny that culpability usually functions as the crucial guide to defensive permissions.¹⁷

Now for the Augmentation Argument. The key to unlocking *Lethal Conditional Steal* can be found in *Modestly Pre-emptive Escalated Steal*. Now *Modestly Pre-emptive Escalated Steal* differs from *Lethal Conditional Steal* in the following respect. In *Lethal Conditional Steal*, Tom immediately issues a conditional threat, to the effect that Daisy either surrenders her money or faces lethal attack. In *Modestly Pre-emptive Escalated Steal*, by contrast, Daisy permissibly adopts lethal defensive measures in the knowledge that Tom *will* escalate to lethal levels if he encounters any defensive resistance. Yet the deep structure of *Modestly Pre-emptive Escalated Steal* and *Lethal Conditional Steal* is similar. Thus we have a route that takes us from the permissibility of the former to the permissibility of the latter.

How so? In *Lethal Conditional Steal*, Tom threatens Daisy with death unless she gives him her money. A defensive response of only H_{STEAL} would be inefficacious against Tom's deadly threat. So what is Daisy permitted to do? I will make three points.

The first of them is this. Like *Modestly Pre-emptive Escalated Steal*, Tom conforms to the profile of the coercive conditional threatener in *Lethal Conditional Steal*. His explicit issue of the conditional lethal threat confirms that he has—and is moreover currently exercising—this conditionally lethal disposition. Daisy can permissibly respond to this threat with defensive violence, just as we saw with *Modestly Pre-emptive Escalated Steal*. But how much violence is she entitled to apply?

The second point—and the answer to the last question—is that Daisy is entitled to represent Tom's lethal conditional threat as a threat which *massively augments* an existing lesser threat of H_{STEAL} . The greater conditional threat facing Daisy (H_{DEATH}) is equivalent to H_{STEAL} , plus whatever further amount of violence (H_{DEATH} minus H_{STEAL}) makes up the difference between them. If Tom threatens Daisy with

¹⁷ For relevant doubts, see, for example, McMahan (2005) and Thomson (1991). Thomson thinks only that the existence of an objectively unjust threat is what matters, while McMahan thinks that responsibility, rather than culpability, is the crucial liability-generating property. These accounts have proven to be highly influential in the literature.

aggressive violence that is *at least* equivalent to H_{STEAL} , then she is permitted to respond defensively to *that* part of the violence. And if the level of violence she is threatened by surpasses H_{STEAL} , then she is also permitted to match this higher level of violence by augmenting H_{STEAL} with whatever further violence is required to match this further, deadlier aggressive threat. (Hence, the Augmentation Argument.)

This leads us to the third point. The aggressive violence contained in Tom's lethal conditional threat— H_{DEATH} , or the sum of H_{STEAL} and (H_{DEATH} minus H_{STEAL})—compresses into one all the additional steps of greater violence that unfold in escalated sequences in *Escalated Steal* and (in a more compressed way) *Modestly Pre-emptive Escalated Steal*. But the compression of these successive escalatory steps in *Lethal Conditional Steal* does not disturb the underlying facts that generated defensive permissions for Daisy in the overtly escalatory cases. Accordingly, if Daisy can permissibly escalate in *Escalated Steal* and *Modestly Pre-emptive Escalated Steal*, she can permissibly defend herself with lethal defensive violence in *Lethal Conditional Steal*.

One final point of clarification. If it is permissible for Daisy to kill Tom in *Lethal Conditional Steal*, is there a danger that we will have to review the verdict of impermissibility on the original case of *Steal*? No. The difference between them is plain enough: in *Steal*, Tom does not threaten Daisy's life, and Daisy would be killing him merely in order to frustrate his theft. That reason is insufficient for lethal defence. In *Lethal Conditional Steal*, by contrast, Tom is conditionally threatening *Daisy's life*: it is not just her money, but her life, that is on the line. That makes a difference, because it allows Daisy to take account of this personal threat *to her* in working out her defensive permissions.¹⁸ She could avoid exposure to Tom's conditional lethal threat, of course, by giving him her money. It may in fact be unwise of her to risk her life by refusing to accede to him,¹⁹ but her refusal to do so and her deployment of lethal violence against him will not be morally impermissible.

8 Conclusion

It is far from obvious how defence theory should handle escalatory cases. I have argued that the defender may permissibly escalate in these cases just as long as they fall into one of two types. In the first type of case, the aggressor poses a coercive conditional threat, in which he leads the escalatory sequence. In the second type of case, the aggressor actively but conditionally threatens the defender by acting on a disposition that matches the disposition of a coercive conditional threatener. I have attempted to establish that conditional threats that seem non-escalatory on the surface have an underlying normative structure that is akin to the structure of escalatory cases.

This essay has also probed the conditions of defence to see how they apply to escalatory cases and conditional threat cases. I have rejected intertemporal versions

¹⁸ This consideration was vital in dispersing the force of the Analogy Argument in Sect. 3.

¹⁹ See, again, the remarks about futility and *Weekend Escalation* in Sect. 4.

of the necessity and proportionality conditions and argued that a theory of defence should accept the existence of defensive permissions in these cases. The upshot of these arguments is that there may be more permissible defensive violence around than many have budgeted for. But as I noted near the start of this discussion: it takes two to have a fight. A theory of defence should not require defenders to bear the costs of aggressors' escalations or conditional threats.

I conclude this discussion with two points. First, some of these arguments may apply to the *Problem of Bloodless Invasion* in Just War Theory. Imagine that one side, A, threatens another side, B, with bloody invasion unless B cedes its territory to A. If B readily acquiesces, then A will refrain from violence. B will lose its political independence, but it will not face further costs. (Assume, unrealistically, that A is a trustworthy threatener: B can be assured that A will indeed refrain from future violence.) Is B permitted to offer violent resistance to A? Traditional Just War Theory answers this question affirmatively, but that permission cannot be taken for granted if the morality of interpersonal defence lies at the foundation of the morality of war, and if defence theory denies violent defensive remedies against conditional lethal threats.

My arguments in Sects. 6 and 7 cannot fail to speak to the Problem of Bloodless Invasion if the Augmentation Argument is on the right lines, and if the reductionist programme for Just War Theory is justifiable.²⁰ In my view, however, war raises moral problems that ordinary interpersonal encounters do not, and so my focus has been deliberately and squarely on private defence. Broader doubts about the reductionist programme in Just War Theory cannot be explored here.

My second and final point concerns the difference between theory and practice. My arguments are not opposed to violence-discouraging measures adopted in the criminal law. I have no quarrel with them. Even if an aggressor *announces* a lethal threat, does he really mean it? Talk is cheap, after all. And even if the defender suspects the aggressor of having the disposition of a coercive lethal threatener, does she really *know* this?²¹ Does the aggressor really have this disposition, and does the defender know that he will maintain it? We should apply caution and epistemic modesty in answering these questions. For that reason, jurisdictions that do not easily permit lethal remedies in this range of cases may actually be justified overall.²²

²⁰ For useful discussions of this problem, see, in addition to Fabre (2014), Frowe (2014), ch. 5, Lazar (2014), Mapel (2019), McMahan (2014), Rodin (2002: 131–139) and Rodin (2014). (This is not an exhaustive list.) Many of these treatments invest simultaneously in both Just War Theory and private defence, precisely because these authors favour a reductionist approach to Just War Theory.

²¹ This is a point I have already touched on: see the conclusion of Sect. 6.

²² Early versions of some of these ideas were presented at the Workshop on Stones versus Lives: Proportionality and Non-Human Value, American University of Beirut, and the Workshop on Heritage in War, University of Oxford, both of which took place under the aegis of the AHRC-funded Heritage in War Project, led by Helen Frowe and Derek Matravers. I am grateful to them and to other members of the audiences on those occasions, especially William Bulow, Bashshar Haydar, Jonathan Parry, Massimo Renzo, Victor Tadros, and Joshua Thomas, for their helpful and searching comments. Later versions of the material received insightful comments from Gustaf Arrhenius, Iwao Hirose, and Shlomi Segall, and members of an audience at the CAMP Seminar in Leeds. Comments by an associate editor and two reviewers also led to a number of improvements in later versions of this essay. My thanks to them all.

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/>.

References

- Fabre, C. 2014. Cosmopolitanism and wars of self-defence. In *The morality of defensive war*, ed. C. Fabre and S. Lazar, 90–114. Oxford: Oxford University Press.
- Finlay, C.J. 2015. *Terrorism and the right to resist: A theory of just revolutionary war*. Cambridge: Cambridge University Press.
- Frowe, H. 2014. *Defensive killing*. Oxford: Oxford University Press.
- Frowe, H. 2015. The role of necessity in liability to defensive harm. In *The ethics of self-defense*, ed. C. Coons and M. Weber, 152–170. Oxford: Oxford University Press.
- Hurka, T. 2005. Proportionality in the morality of war. *Philosophy and Public Affairs* 33(1): 34–66.
- Lazar, S. 2012. Necessity in self-defense and war. *Philosophy and Public Affairs* 40(1): 3–44.
- Lazar, S. 2014. National defence, self-defence, and the problem of political aggression. In *The morality of defensive war*, ed. C. Fabre and S. Lazar, 11–39. Oxford: Oxford University Press.
- Mapel, D. 2019. National defence and political independence. In *The Oxford handbook of the ethics of war*, ed. S. Lazar and H. Frowe, 204–222. Oxford: Oxford University Press.
- McMahan, J. 1994. Innocence, self-defence and killing in war. *Journal of Political Philosophy* 2(3): 193–221.
- McMahan, J. 2005. The basis of moral liability to defensive killing. *Philosophical Issues* 15(1): 386–405.
- McMahan, J. 2014. What rights may be defended by means of war? In *The morality of defensive war*, ed. C. Fabre and S. Lazar, 115–157. Oxford: Oxford University Press.
- McMahan, J. 2015. The limits of self-defense. In *The ethics of self-defense*, ed. C. Coons and M. Weber, 185–210. Oxford: Oxford University Press.
- Oberman, K. 2020. Killing and rescuing: Why necessity must be rethought. *The Philosophical Review* 129(3): 433–463.
- Øverland, G. 2010. Conditional threats. *Journal of Moral Philosophy* 7(3): 334–345.
- Quong, J. 2020. *The morality of defensive force*. Oxford: Oxford University Press.
- Rodin, D. 2002. *War and self-defense*. Oxford: Oxford University Press.
- Rodin, D. 2014. The myth of national defence. In *The morality of defensive war*, ed. C. Fabre and S. Lazar, 69–89. Oxford: Oxford University Press.
- Statman, D. 2008. On the success condition for legitimate self-defense. *Ethics* 118(4): 659–686.
- Thomson, J.J. 1991. Self-defense. *Philosophy and Public Affairs* 20(4): 283–310.

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