

THE BASIS OF MORAL LIABILITY TO DEFENSIVE KILLING

Jeff McMahan
Rutgers University

There may be circumstances in which it is morally justifiable intentionally to kill a person who is morally innocent, threatens no one, rationally wishes not to die, and does not consent to be killed. Although the killing would wrong the victim, it might be justified by the necessity of averting some disaster that would otherwise occur. In other instances of permissible killing, however, the justification appeals to more than consequences. It may appeal to the claim that the person to be killed has acted in such a way that to kill him would neither wrong him nor violate his rights, even if he has not consented to be killed or to be subjected to the risk of being killed. In these cases, I will say that the person is *liable* to be killed. Although I borrow the notion of liability from legal theory, and although much of what I say will be informed by the literature on liability in both tort law and criminal law, my concern in this article is with moral rather than legal liability.

Liability is different from desert. The claim that someone deserves to be killed implies that there is a reason to kill her even if it is possible for no one to be killed; but the claim that someone is liable to be killed has no such implication. Liability is the broader notion: desert implies liability but liability does not imply desert. Thus, if a person deserves to be killed, it follows that he is liable to be killed, but he can be liable to be killed without deserving to be killed. My focus will be on forms of liability that do not involve desert; I will not consider punitive or retributive killing. My focus will instead be primarily on liability to *defensive* killing, though I will also consider whether there can be liability to killing that preserves life but is not strictly defensive because the person to be killed is not the cause of the threat to be averted. Liability, of course, also extends to forms of harmful treatment other than killing, but for simplicity of exposition I will focus on moral liability to be killed. Much of what I will say, however, can be generalized to other forms of harming.

The question I will address is: “What is the basis of liability to killing?” Or, more precisely: “What must a person who does not deserve to die have done to make it the case that he would not be wronged by being killed?” I will examine two widely held accounts of the basis of liability to defensive killing before sketching the outlines of what I think is the most promising account.

The Rights-Based Account

The *Rights-Based Account* holds that what makes a person liable to defensive killing is that he will otherwise violate a right that is sufficiently stringent for killing to be a proportionate means of prevention.¹ According to its most original and influential exponent, Judith Jarvis Thomson, this account offers a justification for self-defensive killing in cases that pose problems for other accounts. It explains, for example, how it can be permissible to kill an “innocent threat,” a moral agent who through voluntary action poses an unjust threat but is not culpable for doing so. Here is an example.

The identical twin brother of a notorious murderer is driving during a stormy night in a remote area when his car breaks down. Unaware that his brother has recently escaped from prison and is thought to be hiding in this same area, he knocks on the door of the nearest house, seeking to phone for help. On opening the door, the armed and frightened resident mistakes the harmless twin for the murderer and lunges at him with a knife.

Assuming that this is the only way to save his own life, may the twin kill the resident? Most people believe that he would be justified in doing so, even if the resident is acting reasonably in the circumstances and thus is fully excused for the threat he poses to the twin. And the Rights-Based Account provides an explanation of why the twin would be justified: because the resident would otherwise violate the twin’s right not to be killed, the resident lacks a right not to be killed by him. In short, the resident makes himself liable to be killed.

Another type of case that is even more problematic for many accounts of self-defense involves a “nonresponsible threat,” a person who poses an unjust threat without being in any way responsible for doing so. Here is Thomson’s example.

A fat man is enjoying a picnic on a cliff directly above the deck on which you are lying with your leg in traction. Suddenly a villain pushes him off the cliff. If he lands on you he will kill you; but he will survive because you will cushion his fall. You cannot move aside but can save yourself by hoisting your sun umbrella and impaling him on it.²

Again, most people believe that you would be justified in killing the fat man. And again the Rights-Based Account offers an explanation, which is that, because the fat man will otherwise violate your right not to be killed, he lacks a right not to be killed by you. He too is liable.

Some have argued that it is not true that the fat man will violate your right if you do not kill him, since only a morally responsible agent can violate a right. Even on Thomson's own account, for a person to have a right is just for others to be morally constrained in certain ways.³ And just as a falling boulder or a charging tiger cannot be morally constrained and therefore cannot violate a right, so also a person cannot be morally constrained except in the exercise of responsible agency. Since the fat man cannot be morally constrained from having his body used by someone else against his will, he does not threaten your right when his body is hurled at you.

This does not, of course, show that the Rights-Based Account is wrong. If it is correct that the fat man would not violate your right, all this shows is that the Rights-Based Account does not, after all, have the intuitively appealing implication that nonresponsible threats may be liable to defensive killing.

Here is another case that challenges the Rights-Based Account. It is familiar from discussions of the Doctrine of Double Effect.

A tactical bomber fighting in a just war has been ordered to bomb a military facility located on the border of the enemy country. He knows that if he bombs the factory, the explosion will kill innocent civilians living just across the border in a neutral country. But this would be a side effect of his action and would be proportionate to the contribution that the destruction of the facility would make to the achievement of the just cause. As he approaches, the civilians learn of his mission. They cannot flee in time but they have access to an anti-aircraft gun.

The traditional question is how the tactical bomber can be justified in bombing the facility when it would not be justifiable for a terror bomber to drop a bomb in the same spot, producing the same effects, with the *intention* of killing the civilians. My question is different. Assuming that the tactical bomber would be morally justified in dropping his bomb, would the civilians be justified in shooting him down in self-defense?

To explain why this case challenges the Rights-Based Account, I need to distinguish two ways of acting against a right. When one is morally justified in doing what another has a right that one not do, one *infringes* her right. When one acts without justification in doing what another has a right that one not do, one *violates* her right.⁴ Because the tactical bomber acts with justification, he will merely infringe the civilians' rights. The usual assumption is that one does not lose one's rights by virtue of morally justified

action. If that is right, the tactical bomber retains his right not to be killed even though he will otherwise infringe the civilians' rights. According to the Rights-Based Account, therefore, the civilians may not kill him in self-defense.

But this is hard to believe—unless, perhaps, his mission is so important that they are morally required to sacrifice themselves for the sake of its success. Can the Rights-Based Account accommodate the common intuition that the civilians may kill the tactical bomber in self-defense? It cannot be claimed that his right is overridden by morally weightier considerations; for the stipulation that his act would be proportionate entails that the failure of his mission, which would be a consequence of their killing him, would be worse from an impartial perspective than their being killed. Perhaps then, contrary to the common assumption, one may lose rights by threatening to infringe rights. If so, then the fact that the tactical bomber will otherwise infringe their right not to be killed means that he lacks a right not to be killed by them—that is, he makes himself liable to be killed by them.

But this too is hard to believe. For if he has lost his right not to be killed by them, it seems that he can have no right of self-defense against them. Yet intuitively it seems that, just as it is permissible for the civilians to kill the tactical bomber in self-defense, it is also permissible for the tactical bomber preemptively to kill the civilians in self-defense. It is, however, hard to see how the Rights-Based Account can explain this apparent symmetry. How can each party have a right to kill the other while at the same time lacking a right not to be killed by the other? How can one be justified in killing someone to prevent him from doing what he has a right to do?

The case of the tactical bomber is a problem not only for the Rights-Based Account but for other accounts as well. I will return to it at the end of the paper. For the moment I will consider another widely held account of the basis of liability to defensive killing.

The Culpability Account

According to Thomson, the violation of a right is a matter of what one person causes to happen to another. It does not require culpability, responsibility, or even agency on the part of the violator. Thus both innocent and nonresponsible threats may violate rights. Although they may be fully excused, they are nevertheless liable to necessary and proportionate defensive action. Yet Thomson acknowledges that some people believe that culpability on the part of the agent is necessary for the violation of a right, and that even more believe that responsible agency is necessary. The latter believe that nonresponsible threats cannot violate rights and thus are not liable to defensive killing, while the former believe that this is true of both nonresponsible *and* innocent threats.

Some theorists have developed accounts of liability to defensive killing that focus on issues of responsibility and culpability for unjust threats. An unjust threat is a threat of harm that would wrong the victim—harm to which the victim is not morally liable. The *Culpability Account* of liability holds that the criterion of liability is culpability for an unjust threat.⁵ There is a narrow version of this account that assigns liability only for culpably *posing* an unjust threat, while a broader version assigns liability for culpable *responsibility* for an unjust threat. According to the broader version, a person may be liable to killing now if that is necessary to avert an unjust threat for which he is responsible through past culpable action, even if he is now no part of the threat. I will assume that the Culpability Account takes this broader form.

It is possible that the Culpability Account is equivalent to a version of the Rights-Based Account that assumes that culpability is necessary for the violation of a right. Like the Rights-Based Account, this account has trouble with the case of the tactical bomber, since it implies that, because the tactical bomber does not act culpably, he does not make himself liable and thus, presumably, may not be attacked in self-defense by the civilians. But I will put these issues aside in order to focus on a different problem for the idea that culpability is the basis of liability. Consider a case of attempted murder.⁶

A villain, fully intending to kill you, points his gun at you and slowly begins to squeeze the trigger. You and he both reasonably believe that the gun is loaded but in fact it is not. You have just enough time to kill him with your concealed derringer before he pulls the trigger.

Most people believe that you are justified in killing him, for he is culpably attempting to kill you and you reasonably believe that he will succeed unless you kill him first. It is not, moreover, obvious that you will wrong him if you kill him. Yet he in fact poses no threat to you. He is not culpably responsible for a threat to your life. He does not, therefore, satisfy the Culpability Account's criterion of liability.

There are several options. One is to retain the Culpability Account's criterion of liability—culpability for an unjust threat—and conjoin it with an objective account of justification. On this view, because objectively the villain poses no threat to you, he is not liable and you cannot claim necessity as a justification. Your killing him would therefore be objectively unjustified, though it would be fully excused because it would be reasonable given your epistemic limitations.

A second option is to retain the Culpability Account's criterion of liability but conjoin it with a subjective account of justification, according to which justification depends not the objective nature of the relation between you and the villain but on what you believe or reasonably believe

about that relation. On this view, the villain is not liable but because you believe, or reasonably believe, that he *is* liable and that killing him is necessary for self-defense, your killing him would be justified.

A third option is to modify the criterion of liability. Intuitively this may seem to be a case in which the villain is liable to defensive killing. Because he is culpably *attempting* to kill you, it does not seem that you would wrong him if you were to kill him, given that you reasonably believe that killing him is necessary for self-defense and that he is culpably responsible for putting you in a situation in which you have to act on that belief. It seems, in short, that the villain's culpability for the attempt on your life is sufficient for liability even though there is in fact no threat for which he is responsible.⁷

Suppose we accept that a person can be liable to be killed by you in self-defense either if he is culpably responsible for an unjust threat to your life or if he is culpably making an unjust attempt on your life. We may still believe that your killing the villain in the case of the culpable attempt is unjustified (though excused). For the villain's liability is not a sufficient justification for killing him if killing him would be pointless because he is in fact harmless. It would, after all, clearly be unjustified for you to kill him if you *knew* that his gun was unloaded, and it is not clear how your lack of morally relevant knowledge can establish a justification where otherwise there would be none. (This is one of the oddities of the subjective account of justification: that it makes mistake of fact a ground of moral justification rather than excuse.)

It seems, however, that there can be cases in which a person can be liable for a culpable attempt on your life and in which your killing him could be justified even when you *know* that he poses no threat to you. These are cases in which the kind of threat the person is attempting to pose arises from a different source and can be averted only by killing the person who is making the attempt. Here is an example.

Aware that a villain plans to kill you, you begin to carry a gun. On one occasion you have the opportunity to empty the bullets from his gun and you do so. Immediately thereafter, he confronts you in an alley and tries to fire. As he continues to pull the trigger in frustration, you see that a second villain is preparing to shoot you from behind a narrow basement window (it is a tough neighborhood). Unable to flee in time and also unable to fire with accuracy through the tiny window, you can save yourself only by shooting the first villain, causing him to slump in front of the window, thereby blocking the second villain's line of fire.

According to the Culpability Account's revised criterion of liability—culpability for an unjust threat *or* culpability for an unjust *attempt*—the first villain is liable to be killed by you because he is making a culpable attempt on your life. You know that this villain is harmless but your killing him

would not be pointless, for it is necessary to prevent yourself from being killed by the second villain. My intuition, which is not shared by everyone, is that you would be justified in killing the first villain. Because he is culpably attempting to pose the same threat you in fact face from another agent, it is difficult to accept that you would wrong him by killing him in order to avert that threat.

Perhaps another case may push us even further.

All of the details are the same as in the previous case except that the second villain is preparing to shoot not you but another innocent person who has appeared in the alley. You can prevent the second villain from killing this innocent person only by killing the first villain, thereby causing him to drop to the ground and block the window.

Again, my intuition is that you would be justified in killing the first villain. My sense that he would not be wronged by being killed by you is slightly weaker in this case than in the previous case, but even in this case I doubt that he has any persuasive ground for complaint.

In short, I suspect that liability to defensive killing can extend beyond culpability for an unjust threat. But once liability becomes detached from threat, various questions arise. One is: "Culpability for what?" These cases suggest, at least to many people, that we should accept an extended criterion of liability to defensive killing that recognizes liability for culpable attempts. But why stop there? It seems arbitrary to draw the line at attempts. Suppose that in the initial case involving the second villain, the first villain follows you into the alley with the intention of killing you but decides to check his gun before firing. Just as the second villain appears at the window, the first villain discovers that his gun is empty. As in the original case, the only way you can prevent the second villain from killing you is to kill the first villain, who not only poses no threat but also now makes no attempt to kill you—though he would kill you if he could. My sense is that his merely discovering that an attempt would be futile cannot make the difference between his being liable and his not being liable.

But is it plausible to suppose that a person can be liable to be killed simply because he *would* unjustly kill an innocent person if he could? And if we accept this as a basis of liability, should we stop here or go further? What about someone who now intends to kill an innocent person in the future? Does it make a difference whether he is actively preparing for the killing? (These questions are crucial to understanding the permissibility of preventive defense.) What about someone who intends to kill an innocent person if certain conditions obtain? Does it matter how likely it is that the conditions will obtain?

Other possibilities are suggested by further variations of the case involving the threatening villain behind the basement window. What if the person

who might be killed in order to block this villain's fire is not now attempting to harm you but made an attempt on your life earlier in the day? Or a month ago? What if, instead, he made an attempt on someone else's life a year ago? What if he is not now attempting to kill anyone but is a murderer who was never punished? Or a murderer who has been punished but remains unrepentant? What if, though harmless, he is simply a thoroughly nasty person?

Clearly we must draw a line before we reach this final possibility. I find it plausible that the person who is culpably but futilely attempting to kill you would not, in our first three cases, be wronged if you were to kill him. But once liability is divorced from culpability for an actual threat, I do not know where to draw a principled line at which culpability ceases to be a basis of liability.

There is another problem. If, as the Culpability Account originally asserted, liability to defensive killing arises only from culpability for an unjust threat, that implicitly limits the purpose for which the liable person may be killed to the elimination of the threat. But if liability is divorced from responsibility for an actual threat, it is an open question for what purposes a person who is liable to be killed may justifiably be killed. This problem is illustrated by the case in which the purpose of killing the first villain is not to defend the person he is attempting to kill but to defend a different person from a different threat. But there are other possibilities. We might, for example, kill those who are liable for culpable attempts in order to use their organs for transplantation. It is unclear, however, whether there is a stable, principled line between this unappealing suggestion and the view that it is permissible to kill the first villain to prevent an innocent bystander from being killed by the second villain.

The cases I have considered thus far pose the question whether certain forms of culpability are sufficient for liability. But one point on which the Culpability Account is firm is that culpability is *necessary* for liability. It thus denies that either innocent or nonresponsible threats can be liable to defensive killing. Many people will regard this as a sufficient basis for rejecting the account. For most people agree with Thomson that it can be permissible to kill innocent and nonresponsible threats in self-defense.

I disagree. I think that there is no moral basis for holding nonresponsible threats liable.⁸ And it is possible that there is no basis for holding certain innocent threats liable. But *some* innocent threats *are* liable. In other words, culpability is not necessary for liability. Here is an example of liability without culpability.

A person keeps his car well maintained and always drives cautiously and alertly. On one occasion, however, freak circumstances cause the car to go out of control. It has veered in the direction of a pedestrian whom it will kill unless she blows it up by using one of the explosive devices with which pedestrians in philosophical examples are typically equipped.

This conscientious driver is an example of one type of innocent threat, which I call an “inadvertent threat.” The pedestrian has done nothing to make herself liable to be killed but the driver will kill her unless he is himself killed first. Yet he is in no way culpable. The unjust threat he poses is not the result of wrongful intent, recklessness, or negligence; it is the result of sheer bad luck. On what basis, then, may he be said to be liable?

What makes him liable is that, as a morally responsible agent, he voluntarily chose to set a couple of tons of steel rolling as a means of pursuing his ends, knowing that this would involve a tiny risk that he would lose control of this dangerous object that he had set in motion, thereby imperiling the lives of the innocent. It is important to notice that the claim here is not that he is liable because he alone had the chance to avoid the tragic conflict between himself and the pedestrian. For the pedestrian could have avoided it as well by simply staying at home. The claim is, rather, that the driver is liable because he voluntarily engaged in a risk-imposing activity and is responsible for the consequences when the risks he imposed eventuate in harms.

Admittedly, what I am citing as the relevant difference between the driver and the pedestrian is of comparatively slight moral significance. But given that the costs of the driver’s voluntarily engaging in a risk-imposing activity cannot be divided, it is fair that he should suffer them all rather than impose them all on the pedestrian. Yet because he is not culpable for the threat of unjust harm he poses, it might be reasonable to demand that the pedestrian share the costs with him if that were possible. If, for example, the pedestrian could defend her life in a way that would inflict nonlethal harm on the driver but also allow his action to inflict a nonlethal injury on her, she might, particularly if the harm she would suffer would not be severe, be obliged to suffer this harm rather than kill the driver. To kill him when this alternative is available might indeed be to wrong him. (Note that this assumes that liability and proportionality in self-defense are interrelated.)

The Responsibility Account

The case of the conscientious driver suggests an alternative account of the basis of liability. According to what I will call the *Responsibility Account*, the criterion of liability to defensive killing is moral responsibility, through action that lacks objective justification, for a threat of unjust harm to others, where a harm is unjust if it is one to which the victim is not liable and to which she has not consented. For convenience, I will sometimes abbreviate this to “responsibility for an unjust threat.” The conscientious driver satisfies this criterion because he is a morally responsible agent who has acted voluntarily in a way that foreseeably imposed risks on others and now threatens the pedestrian with unjust harm. And he acts without

objective justification because, although he may act reasonably given his epistemic limitations, it is not justifiable to drive a car that is in fact going to go out of control. Yet he is in no way culpable for the threat he poses. The epistemic constraints under which he acts are fully exculpating.

The Responsibility Account treats liability as a matter of degree. Responsibility for an unjust threat can vary in degree, and liability varies concomitantly. A person's liability is therefore greater when his action is culpable, and the degree of his liability varies with the degree of his culpability. The degree of a person's liability is manifest in the severity of what may be done to him without wronging him—that is, in the stringency of the proportionality requirement.

It may be instructive to compare liability to defensive action to liability in tort law. Just as we may think of liability in torts as a matter of corrective justice, or justice in the distribution of harm *ex post*, so we may think of liability to defensive action as a matter of preventive justice, or justice in the distribution of harm *ex ante*. According to the Culpability Account, liability to defensive action corresponds to fault liability in torts. But the Responsibility Account treats liability to defensive action as strict, as in the case of the conscientious driver. It says to all people who engage in socially permitted risk-imposing activities such as driving: “You know that if you drive you impose a very small risk on other innocent people. If you choose to drive, the consequences are your responsibility unless others also contribute to the outcome through their own risk-imposing activities. You will be liable to defensive action even if you satisfy the relevant standards of due care.”

It is, perhaps, significant that there is no account of liability in torts that corresponds to the Rights-Based Account as it is usually understood, for example by Thomson. Suppose the fat man who is pushed off a cliff actually lands on you but merely breaks your leg rather than killing you. There is no conception of strict liability that is strict enough to hold him liable for damages.

It is also true, however, that if the conscientious driver were merely to injure the pedestrian rather than kill her, he too would escape liability in tort law. His act would not be among those that tort law subjects to a standard of strict liability. More generally, the range of injurious acts for which tort law holds people strictly liable is narrower than the range of threatening acts for which the Responsibility Account holds people liable to defensive action. I believe, however, that tort law would correspond more closely to our sense of corrective justice if it were to hold the conscientious driver liable. Yet the provisions of tort law have to be formulated to take account not only of considerations of justice but also of consequences. If people like the conscientious driver were held strictly liable *ex post*, that might help to deter people from driving for frivolous reasons, which would be desirable, but it might also have undesirable effects as well. I cannot pursue these issues here.⁹

I believe that some version of the Responsibility Account will turn out to be the best account of liability to defensive killing. Like the other accounts, however, the Responsibility Account, at least in the form in which I have stated it, faces a battery of objections. In the remainder of this article I will review some of these objections and suggest some initial responses.

Suppose that the only way you can prevent yourself from being killed by a culpable attacker is to kill his mother. If you do kill her, can you then claim that she was liable to be killed because, as a morally responsible agent, she voluntarily chose to engage in an activity (having a child) that had a tiny probability of resulting in an unjust threat and that this made her responsible for the threat you faced from her son?¹⁰ Obviously not. But it is less obvious what the right explanation is of why the mother is not liable.

One explanation is that the woman's act of having a child is causally too remote for her to be responsible for the threat he poses. It might be argued that just as tort law insists that a person must have been the "proximate" cause of harm to another in order to be liable to repair that harm, so a person must be the proximate cause of an unjust threat to another in order to be liable to defensive killing. This seems, however, to be too strong a requirement. Suppose that a wholly innocent person is about to be tortured as a terrorist suspect under newly adopted rules that wrongfully permit certain forms of torture during interrogation. She cannot defend herself against the person who will actually conduct the torture but, through a remarkable mechanism known only to those who devise philosophical examples, she can prevent the torture by inflicting a proportionate harm on an official who provided sophisticated arguments for the legal permissibility of the new rules, thereby making it possible to implement them. This official seems responsible for the practice of torture in a way that makes him liable to defensive action. He would not be wronged if this woman were to inflict a proportionate harm on him as a necessary means of avoiding torture that she would not have faced were it not for his action. Yet the action that makes him liable is causally remote from the unjust threat the woman faces. The torturer, not the official, is the proximate cause of that threat. Still, it seems that the causal connections are of the right sort for the transmission of moral responsibility. I will not try to determine what sorts of causal connection are the right sorts. But it seems that the official's causal role in implementing the policy of torture is part of the cause of the torturer's action and that this makes him responsible to a significant degree for what the torturer does. By contrast, the mother's action in causing her child to exist does not in anything like the same way cause his subsequent acts and does not make her to any degree responsible for them.

Not only are there some who, like the official, are liable for an unjust threat without being its proximate cause; there are also some who are the proximate cause of an unjust threat but are not liable. This is true, for

example, of the falling fat man. But it also seems true of some innocent threats who might seem to be liable under the Responsibility Account as I have stated it. In earlier work, I assumed that the only instances in which a person can be altogether lacking in responsibility for harm he causes are those in which he causes harm independently of his responsible agency. I therefore assumed that whenever a morally responsible agent causes harm through his own voluntary action, he is to some degree responsible for it, even if he is fully excused for causing it because of some excusing condition or simple bad luck.¹¹ But I now believe that this was a mistake. Here is an example that shows why.

A villain has secretly tampered with a man's cell phone in such a way that if the man presses the "send" button, he will detonate a bomb to which you have been tied by this same villain. The villain has trussed you up so that you cannot escape or alert the man with the cell phone to your plight. But the villain has given you a weapon with which you can kill the man with the cell phone.

Suppose the cell phone operator has pressed in a number and is about to press the send button. He now inadvertently poses an unjust threat to you—that is, a threat to which you are in no way liable. He is, moreover, a morally responsible agent whose pressing the button would be a free, voluntary act. But, even though he is the proximate cause of the threat to you, he seems in no way responsible for that threat. Consequently, he seems no more liable than the falling fat man. It is possible that you could be justified in killing him, but you could not plausibly claim that he had made himself liable to be killed.

In our earlier example, the killer's mother could reasonably have foreseen that there was a very small probability that if she had a child he would become a murderer. But in that case there was not the right sort of causal connection between her having a child and that child's later posing an unjust threat to transmit responsibility for that threat back to her. In the example of the cell phone operator, there is certainly the right sort of causal connection between the man's action and the unjust threat you face. The necessary ingredient of responsibility that is missing in this case is foreseeability.

There is, of course, a problem about how probable a threat or harm must be in order to count as foreseeable. I cannot address that problem here; for present purposes I will sidestep it by simply invoking the notion of a "risk-imposing activity"—that is, an activity of a type that is known or can be seen to impose a risk of significant harm on others, even if the probability is comparatively low, as it is in the case of driving. I will assume that it is a condition of responsibility for an unjust threat that the action that gave rise to the threat either was of a risk-imposing type or was such that in the circumstances the agent ought to have foreseen that it carried a non-negligible risk of causing a significant unjust harm.

While the conscientious driver and the cell phone operator are both the inadvertent but proximate cause of an unjust threat, only the driver is engaged in a risk-imposing activity. Pressing a button on a cell phone is not the kind of act that typically carries any risk to others, however slight. This is why the cell phone operator cannot be held responsible or therefore liable for the harm his pressing the button would cause. (Having a child is also not a risk-imposing activity, but this is not because it is not a type of activity that leads ultimately to risks to others. It is because any subsequent risks or harms cannot be traced in the right way back to the activity itself.)

In short, the cell phone operator is not liable to defensive killing according to the Responsibility Account because he fails to meet one of the conditions of moral responsibility—the foreseeability condition—for the unjust threat he poses. Although he is a morally responsible agent who acts voluntarily and without objective justification (since his pressing the send button is objectively unjustified), he is nevertheless, like the falling fat man, a nonresponsible threat.

I find it a virtue of the Responsibility Account that it identifies a basis for attributing liability to the conscientious driver. Intuitively, however, it seems to matter what his reason is for driving. On considering the example, we naturally assume that this person is driving for one of the self-interested reasons for which people typically drive: to go to the store, to get to work, or whatever. But some people, such as ambulance drivers, drive for more important reasons—reasons for which other people reasonably want them to drive despite the risks they impose. If we imagine a variant of the case of the conscientious driver involving an ambulance driver on the way to the scene of an accident, we are likely to question whether his voluntarily engaging in this risk-imposing activity can make him liable to defensive killing by the pedestrian when the ambulance goes out of control.

It is tempting to suppose that the ambulance driver's action is justified in a more robust sense than the action of the conscientious person who drives for reasons of self-interest and that this blocks an attribution to him of liability. But objectively neither driver is justified. For in each case, however faultless or even meritorious the driver may be, there is an objective description of what he will do unless he is killed—killing an innocent pedestrian—under which it is clear that his action is unjustified. (I am assuming that getting to the accident site is not sufficiently important to make it justifiable for the ambulance driver knowingly to kill a pedestrian if that were necessary to get there.)

This, of course, presupposes an objective account of justification. But even a subjective account of justification is unlikely to be able to distinguish between the two drivers in a way that supports an attribution of liability to the self-interested driver but denies that attribution to the ambulance driver. For any plausible subjective account that implies that the ambulance driver's action is justified will have the same implication for the conscientious

driver's self-interested action, given how low the antecedent subjective probability of his car's going out of control was.

Suppose, however, that we were to find a plausible basis for claiming that, while the conscientious driver is liable for having chosen to engage in a risk-imposing activity, the ambulance driver's *reason* for engaging in that same activity shields him from liability in the same circumstances, even if his action is objectively unjustified. We would still face a line-drawing problem. For people's reasons for engaging in a risk-imposing activity such as driving lie along a spectrum from culpable to morally neutral to praiseworthy. If we claim that the ambulance driver's reason for driving blocks an attribution of liability to him, what should we say about a person who drives to take his child to the doctor? Or to take his sick dog to the vet? Or to present a paper on liability to defensive killing at a philosophy colloquium? There does not seem to be any point at which a principled line can be drawn between liable and not liable on the spectrum of cases between the conscientious driver and the ambulance driver. (And even if we could find such a point, identifying it would be largely irrelevant to matters of practice, since in cases of self-defense, and particularly cases of defense against inadvertent threats, agents tend to have little or no reliable information about the reasons for which those who pose an unjust threat have acted.)

It may seem that the ambulance driver is similar to the tactical bomber in that both unintentionally threaten the innocent as a side effect of engaging in morally desirable action. According to the Responsibility Account, however, they are quite different, for the tactical bomber acts with objective justification while the ambulance driver does not. This fact about the tactical bomber may, indeed, enable the Responsibility Account to deal in an intuitively satisfactory way with the problem the tactical bomber poses.

In the case of the tactical bomber, the civilians are not liable. The tactical bomber will therefore wrong them if he drops his bomb. Yet he acts with objective moral justification. And, according to the Responsibility Account, justification defeats liability: one is liable to defensive killing by virtue of responsibility for an unjust threat only if one acts without justification. This seems intuitively plausible; for it is hard to see how one's moral immunity to being killed could be compromised merely by one's acting in a way that is morally justified.

As I noted earlier, however, it seems intuitively that the civilians are permitted to kill the tactical bomber in self-defense. How can that be if he is not liable? I suggest that people are permitted a necessary and proportionate defense against being wronged either by unjustified *or justified* action. People are not morally required, or at least not always, simply to submit to being wronged by another's morally justified action. They are, at least sometimes, morally permitted to resist. This is a special form of justification for killing; it does not appeal either to liability on the part of the victim or to the overriding moral importance of avoiding disastrous consequences.

It is essential to this form of justification that it applies only to those who would otherwise be wronged by morally justified action. A person who acts with justification to threaten another with harm to which the other is morally liable does not threaten to wrong that other person. There can be no permissible defense against harm to which one is morally liable.

Because the tactical bomber's justified action would wrong the civilians, they are permitted a proportionate defense. And killing him would be proportionate. But because he is not morally liable to be killed by them, they will wrong him if they kill him in self-defense. He too, therefore, is permitted a proportionate defense against their permissible defensive action. He is justified in killing them in preemptive self-defense.

This reasoning establishes the intuitive moral parity between the tactical bomber and the civilians. Neither party to the conflict is liable to be killed by the other, but each is justified in killing the other in self-defense.

This reasoning is also compatible with the plausible view that third parties are not, other things being equal, permitted to intervene. Although the fact that each party threatens to wrong the other seems to provide a ground for third party intervention, that is canceled by the fact that each party also acts with justification.

This understanding of the case of the tactical bomber can, it seems, also be stated in the language of rights. The civilians have a right not to be killed. Because the tactical bomber acts with justification, his killing them would not violate but would infringe their rights. But people are permitted a necessary and proportionate defense of their rights against both violation *and infringement* (though perhaps what is proportionate to prevent violation may be disproportionate as a means of preventing infringement). Hence the civilians are permitted to kill the tactical bomber in self-defense. In doing so, however, they would infringe his right not to be killed by them. For while it is true that one may lose rights by threatening to violate rights, one does not lose rights by threatening to infringe rights. For one cannot lose rights merely by acting in a way that is morally justified. Thus, because the tactical bomber would merely infringe the rights of the civilians by killing them, he retains his own right not to be killed by them. It is therefore permissible for him to defend his right against infringement by killing them in preemptive self-defense. In short, he is justified in dropping his bomb, though that would infringe the rights of the civilians; they are justified in killing him to prevent him from infringing their rights; but their killing him would infringe his rights; hence he is also justified in killing them to prevent them from infringing his right not to be killed. Moreover, disinterested third parties may not intervene on behalf of either party to the conflict; for while they may, and sometimes must, prevent the *violation* of a right, they may not prevent the *infringement* of a right, since they must not kill a person in order to prevent him from engaging in morally justified action.

Only those whose rights would be infringed may do that; others must remain neutral.

The ease of translation here suggests the possibility of formulating a variant of the Rights-Based Account that is equivalent, or at least extensionally equivalent, to the Responsibility Account. We may, perhaps, think of the Rights-Based Account as developed by Thomson as only one way of filling out a formal structure that may yield a rather different theory when filled out with substantive assumptions different from those made by Thomson. Among these assumptions are that the conditions for the violation or infringement of a right are the same as the conditions of moral responsibility for an unjust harm and that it is permissible to resist the infringement as well as the violation of one's rights (though the proportionality restriction may be more stringent in the former case than in the latter). But while it is possible that all that I say may be translatable without loss into the idiom of rights, I will not investigate this possibility further here.

One worry that may arise is that the explanation I have given of why the civilians are justified in killing the tactical bomber in self-defense also applies to the potential victims of the falling fat man and the cell phone operator. If the civilians may kill the tactical bomber to prevent him from wronging them, why is it not similarly permissible for you to kill the falling fat man and the cell phone operator in order to prevent them from wronging you? Admittedly, there are many civilians and only one tactical bomber, whereas the exchange of lives in the other two cases would be one for one. But this consideration is offset by the fact that the tactical bomber, but not the fat man or the cell phone operator, poses an unjust threat through action that is objectively justified. In short, while the tactical bomber will wrong his victims with justification, both the fat man and the cell phone operator will wrong you *without justification*. Why may you not kill them?

The answer is that they will not in fact wrong you. Whether you are *wronged* is not a matter only of what happens to you. It also depends on facts about agency and responsibility. A falling boulder or a charging tiger may harm you but not wrong you. Moral responsibility is a precondition of wronging, and neither the fat man nor the cell phone operator satisfies the conditions of moral responsibility for the unjust threat he poses. The threat posed by the fat man is not a product of his responsible agency, while the threat posed by the cell phone operator cannot be foreseen by him—that is, he cannot know that he poses that threat. (Expressed in the language of rights, the claim here is that neither the fat man nor the cell phone operator will either violate or infringe your right not to be killed. Both violation and infringement of rights require moral responsibility. Indeed, infringement, as I understand it, refers to action that is morally justified. And the fat man, for example, is not justified in falling on you. Involuntarily falling is neither justified nor unjustified.)

In short, the tactical bomber, the fat man, and the cell phone operator all fail to meet the Responsibility Account's conditions of liability. But the tactical bomber does meet the account's conditions of responsibility; thus, his action will wrong the civilians. The reason he is released from liability is that he acts with justification. The fat man and the cell phone operator, by contrast, are not released from liability because they act with justification. The reason they are not liable is that they fail to meet the conditions of moral responsibility; and this means that though they act without justification, they will not wrong you even if they kill you.

Because they are not liable, all three—the tactical bomber, the fat man, and the cell phone operator—would be wronged by being killed. But at least in the circumstances as described, only the tactical bomber may permissibly be wronged. Acting with justification does not, it seems, afford the same degree of immunity as lack of moral responsibility.

Recall now three other cases we have discussed. The resident in the case of the identical twin, the conscientious driver, and the ambulance driver all seem to meet the conditions of responsibility for the unjust threats they pose. All will violate their victim's right not to be killed. For none acts with objective justification. It seems that all three are liable to defensive killing according to the Responsibility Account. (Perhaps implausibly in the view of most people, the one who may come closest to avoiding liability, according to the Responsibility Account, is the resident. For we can imagine that story in ways in which it would be virtually impossible for him to know, in the time in which it is reasonable for him to act, that the threat he poses is unjust. Even so, we may think it a relevant difference between him and the cell phone operator that he at least knows that he poses a threat, and thus is responsible for ensuring that it is not an unjust threat.)

It is, perhaps, tempting to think that these last three cases should be treated in the way I have suggested that we treat the case of the tactical bomber. Consider, for example, the conscientious driver. I have argued that he is liable by virtue of having chosen to engage in a risk-imposing activity. But this, it might be argued, is too flimsy a basis for a claim of liability. For many of the normal activities of life are risk-imposing, and it would be absurd to assign liability to people for simply living their lives in normal ways. Even just going out in public imposes risks of serious harm on others. One may, for example, be an asymptomatic carrier of a flu virus, which might prove fatal to a very young or very elderly person whom one might expose to it. (This, unlike the possibility of one's cell phone being secretly transformed into a detonator, is a possibility of which one is, or ought to be, aware.)

This suggestion is likely to be urged by a proponent of the Culpability Account, who claims that fault is necessary for liability. To be liable, on this view, one must depart in certain specific ways from the norms of ordinary social life. The defender of this view could argue, moreover, that the denial

that the conscientious driver is liable does not entail that the pedestrian must allow herself to be killed—that is, that she may not kill in self-defense. For my own discussion of the tactical bomber concedes that there are justifications for killing that appeal to neither the victim's liability nor the necessity of averting a catastrophe. Why cannot we say that the pedestrian may kill the conscientious driver in self-defense simply because the conscientious driver will otherwise wrong her, or violate her rights, and that the conscientious driver, being free of liability because he is not culpable, is likewise permitted to kill the pedestrian in preemptive self-defense?

The problem is that this sort of resolution does not seem to be available in the conflict between the conscientious driver and the pedestrian. The case of the tactical bomber is special: what absolves him of liability is that he acts with objective justification. But this is not true of the conscientious driver. Even if he acts reasonably, he has no positive moral justification for doing what will kill the pedestrian. Since he will therefore *violate* the pedestrian's right not to be killed, it is unclear how he could be free of liability. Is it because he fails to meet the conditions of moral responsibility? No, because he clearly does meet those conditions. There is, moreover, a relevant difference between him and the pedestrian, which is that he chose to act in a way that exposed the pedestrian to a risk of serious harm whereas the pedestrian imposes no comparable risk on him. I believe, therefore, that it would be unjust if the conscientious driver, having chosen to do what exposed the pedestrian to risk, were then to kill her preemptively to prevent her from defending herself from the harm that he would otherwise, and without justification, cause her.

I continue to believe, therefore, that the Responsibility Account's implications for the case of the conscientious driver are more plausible than those of the Culpability Account. I do not, however, point to this in triumph, for this is not the kind of case against which we can usefully test a theory's implications. It is a case about which most of us have only weak or doubtful intuitions, and thus is precisely the sort of case for which we need guidance from a theory.

In criticizing the Culpability Account, I suggested that it seems plausible to regard those engaged in culpable attempts as liable to defensive action, even though they would otherwise cause no harm. If this is an objection to the Culpability Account, it is equally an objection to the Responsibility Account, which also insists on causation as a condition of liability. (The problem of attempts arises only in cases in which the Responsibility Account coincides with the Culpability Account—namely, cases involving culpable responsibility for an attempt to cause an unjust harm.)

In my earlier discussion of culpable attempts, I left the line-drawing problem unresolved. I will conclude by offering a brief argument in favor of retaining causation as a condition of liability. The argument appeals, not to attempts, but to a variant of the case of the conscientious driver. Suppose

that when the driver's car goes out of control, it threatens another driver rather than a pedestrian. In this case, the imposition of risk has been reciprocal. Both drivers made the same choice to expose others to risk. But only one, through moral bad luck, now threatens unjust harm to the other. Most people are willing to allow that moral luck may affect liability and thus believe that only the driver who poses an unjust threat is liable. Most people, I think, would agree that the driver who is threatened would be justified in acting in self-defense but that the threatening driver would not be justified in acting to preempt the other's defensive action. If that is right, the mere imposition of risk—or, in the case of attempts, mere culpability—is not sufficient for liability in the absence of causation. Causal responsibility for an unjust threat is a necessary condition of liability to defensive killing.¹²

Notes

1. This view is developed and defended in Judith Jarvis Thomson's important article, "Self-Defense," *Philosophy and Public Affairs* 20 (1991): 283–310. Thomson uses the term "innocent aggressor" to refer to what I call an innocent threat. What she refers to as an innocent threat is what I call a nonresponsible threat.
2. Thomson, p. 287. In the story as she tells it, you can use your awning to deflect him past the edge of the deck onto the road below. Since there are accounts of the distinction between killing and letting die that would classify your deflecting him as allowing him to die rather than killing him, I have altered the example in a way that makes your act an uncontroversial instance of killing.
3. Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), p. 77.
4. I draw this distinction differently from the way Thomson does. *Ibid.*, p. 122.
5. One question I will not pursue is whether a person may be liable for being culpably responsible for a threat that is not unjust because the victim is liable to that threat. There are two relevant types of case. In one, a case of "unknown justification," an agent poses a threat to which the potential victim is in fact liable though the agent acts on a wrongful reason in ignorance of the reasons that could justify his action. In the other, a case of "ostensible justification," an agent poses a threat to which the potential victim is liable but does so for a wrongful reason on the pretense of acting on the basis of the justifying reason. In these cases, it is controversial not only whether the agent is justified but also whether the potential victim has a right of self-defense.
6. Some parts of this section draw on my "Self-Defense and Culpability," forthcoming in *Law and Philosophy*. That paper comments on an account of the right of self-defense advanced by Kimberly Kessler Ferzan in "Justifying Self-Defense," which is also forthcoming in *Law and Philosophy* and from which I have learned a great deal.
7. This is the view defended by Ferzan in "Justifying Self-Defense."

8. I argue for this in *The Ethics of Killing: Problems at the Margins of Life* (New York: Oxford University Press, 2002), pp. 398–411.
9. For relevant discussion, see Steven Shavell, “Strict Liability versus Negligence,” *Journal of Legal Studies* 9 (1980): 1–25; and Douglas Husak, “Vehicles and Crashes: Why is this Moral Issue Overlooked?,” *Social Theory and Practice* 30 (2004): 351–70.
10. I owe this objection to Kai Draper.
11. See *The Ethics of Killing*, pp. 401–5.
12. For perceptive written comments on an earlier draft, I am grateful to Shawn Bayern, Ruth Chang, Kimberly Kessler Ferzan, Peter Vallentyne, and Alec Walen. I am also indebted to Judith Jarvis Thomson for comments on an earlier manuscript from which this article draws. For exceptionally stimulating discussion, I am grateful to Larry Alexander, Allen Buchanan, Agnieszka Jaworska, Sanford Kadish, Shelly Kagan, Frances Kamm, Christopher Kutz, William McMahan, Eric Rakowski, Larry Temkin, Eric Wampler, and, especially, Samuel Scheffler.