

I **Self-defence**

Many just war theorists have sought to explain the permissibility of killing in war by pointing out that there are other exceptions to the general prohibition on taking life. Most notable among these is the exception that permits the killing of an attacker in self-defence or in defence of other people. When our lives are threatened, most of us think that it is permissible to try to defend ourselves even if we can do so only by inflicting very great or even lethal harm upon our attacker. We also think that someone who defends us against such an attack, even if they kill our attacker, is pretty heroic. So if a case can be made that killing in war is a form of defensive killing, we may have the beginnings of an explanation of why killing in war is permissible.

The idea of war as self-defence has had a profound influence on just war thought across the centuries. But the relationship between war and self-defence is less than straightforward, not least because there is no general consensus amongst either philosophers or legal theorists about when or why self-defence is permissible. Nor is there consensus about how strongly we should interpret the relationship between war and self-defence. However, since many just war theorists grant the existence of at least some relationship between war and self-defence, it will be worth our while to look at some of the main tenets of influential accounts of permissible self-defence. Several of the concepts that we use in just war theory, such as the requirements of necessity and proportionality, certainly seem to have their roots in thoughts about self-defence, even if they are to be understood differently when they are used in the context of war. Once we have familiarised ourselves with some important aspects of self-defence, we can look in more detail at the relationship between war and self-defence.

Proportionality and necessity

Some people believe that it is never morally permissible to use force against another person. This position is usually described as pacifism. A committed pacifist will believe that even when one's life is threatened by a murderous attacker, morality does not permit one to use force to prevent that attack. Nor may one use force to defend other people's lives. Most people are not pacifists. Most people think that, under certain circumstances, it can be permissible to use force, including lethal

force, against another person. The sorts of circumstance that people making this claim usually have in mind are those in which the force will save either one's own life or someone else's life. But the permission of self-defence is not unlimited. For example, most people think it permissible to use force only when doing so is necessary. And most people think that there are limits to when I can use even necessary force to save my own life. I probably shouldn't kill you and eat you even if doing so is necessary to ward off the starvation that will otherwise kill me. The purpose of philosophical accounts of self-defence, then, is to explain when and why using defensive force is permissible.

Proportionality

That defensive force must be proportionate is a standard constraint on permissible defence. As we shall see, judging proportionality in the context of war has proved both complex and controversial. For our purposes here, however, we can stick to a fairly basic account of proportionality. Roughly, the proportionality condition stipulates that the harm I inflict upon my attacker must not significantly outweigh the good that I hope to thereby secure. Usually, the good that I hope to secure is the averting or mitigating of the harm that they threaten to inflict upon me. While I can permissibly break their leg to stop them from breaking my leg, I cannot break their leg to stop them from pinching me.

How can we tell if a defensive harm is proportionate to a threatened harm? When the two harms are equal, this seems easy enough. If I kill you to stop you from killing me, the good that I secure (preventing my death) is equal to the defensive harm that I inflict (your death). Killing you thus seems proportionate. But most people think that it is not only when I face a lethal threat that I may lethally defend myself. I may also lethally defend myself against non-lethal harms such as rape, torture, paralysis or kidnap.

There seem to be two possible explanations of how such defensive killings could meet the proportionality requirement. The first is to suggest that while these harms of rape, torture and so on are not the same as death, they are nonetheless of the same *magnitude*. We might argue that being raped or tortured is as bad as being killed, and thus these harms are in fact equal to death. This would explain why I can kill a person who will otherwise rape or torture me without falling foul of the proportionality condition.

But even if we can make this case for rape or torture, it seems somewhat less plausible with respect to paralysis or kidnap. Being paralysed is certainly a very serious harm. But it does not seem as bad as being killed, at least when the paralysis is only partial. Similarly, kidnapping does not in itself seem to be a harm comparable to death. It may be terrifying and a gross invasion of one's liberty, but in the absence of additional harms such as torture or rape, being kidnapped does not seem to be a harm of the same magnitude as being killed. Of course, we could simply deny that it is proportionate to use lethal force to prevent these non-lethal harms. But if we think that people are allowed to lethally defend themselves

against at least some lesser harms, we will need to adopt an alternative account of proportionality that does not insist that the degree of defensive force should be no greater than the threatened harm.

An alternative strategy, then, is to argue for a threshold approach to proportionality. This approach does not require that one should be threatened with a harm equal to death before lethal defence is proportionate, but only that one should be faced with a sufficiently serious harm. Since rape, torture, paralysis and kidnap are all very serious harms, perhaps this is enough to make it permissible to lethally defend oneself against them. On this understanding of proportionality, we have a sort of ranking of harms, whereby one can inflict a defensive harm that is greater than the threatened harm, but still comparable to the threatened harm. This would explain why I can break an attacker's leg not only if he will otherwise break my leg, but also if he will otherwise break my arm or my ankle. A broken leg seems undeniably worse than a broken ankle, but the two harms are nonetheless comparable, such that breaking the assailant's leg to protect my ankle does not seem disproportionate.

Necessity

The necessity condition is also a standard requirement of permissible defence. This condition forbids using more force than one has to in the course of defending oneself. If I can prevent you from killing me either by stamping on your foot or by shooting you in the chest, morality requires that I choose the least harmful means and stamp on your foot. Note, however, that shooting you in the chest would not be disproportionate, since you do pose a lethal threat to me. Also note that the mere fact that a harm is necessary cannot make it proportionate. Even if breaking your leg is the only way in which I can stop you from pinching me, it is impermissible to break your leg. In such a case, I must submit to the infliction of the minor harm rather than inflict a serious harm in self-defence.

Some people describe the necessity condition as a requirement of 'last resort'. One should use force only if one has no other option. Several accounts of self-defence thus endorse the idea that a person must try to retreat from their attacker rather than use force. If I can flee without exposing myself to any great risk, I ought to do so. A notable exception to this, however, might be the issue of retreating within one's own property.¹ Some legal systems do not require a person to leave their home rather than confront an attacker, even if they can safely do so. This reflects many people's intuition that it would be permissible to stay and forcefully defend oneself against an attacker in one's own home, although people differ on how much force one could use rather than flee. If, in order to stay, one will have to use lethal force against an attacker, we might think that morality requires one to flee if it is safe to do so.

People also disagree about the permissibility of using significant force to avert an intrusion into one's home if the intrusion is not accompanied by a threat of physical harm. The 1999 case of Tony Martin – the British man who

shot two would-be burglars, killing one and wounding the other – sparked an enormous debate over the defensive rights of homeowners. Part of the controversy comes from whether it can ever be proportionate to inflict serious physical harm to protect against some non-physical harm such as theft. Some people think that while I am permitted to inflict minor harm – to punch you or shove you – to stop you from stealing my purse, I cannot stab you to stop you from stealing it, no matter how much money is in it. Property alone is not usually thought to be of sufficient importance to make serious harm a proportionate means of protection. But we might think that our homes are more than mere property: that they are ‘our part of the world’, in which we have the right to feel safe. On this view, invading our homes threatens more than just our property rights. This might explain some people’s intuition that it is permissible to inflict serious or lethal harm in order to stay in our homes, even if we could safely retreat from them.

Of course, that force is necessary and proportionate is only part of what makes the use of force an act of permissible defence. As I suggested earlier, killing and eating a person to ward off starvation could be described as a necessary and proportionate force if it is the only way in which I can save my life. However, it hardly seems like permissible self-defence. So what we must investigate now are the other conditions that must be met if defence is to be permissible.

The culpability account

Imagine the following scenario.²

Truck

You are standing in a meadow, innocently minding your own business. Villain, who hates you, is driving towards you in a truck, intent upon running you down. There is nowhere to run. Luckily, you have your antitank gun with you. You can blow up the truck, killing Villain, to save your own life.

Most people think it permissible for you to blow up the truck even though doing so will kill Villain. And an intuitive explanation of this permissibility is that Villain is, after all, a villain. He is maliciously trying to kill you, an innocent person. On the *culpability account* of self-defence, it is Villain’s malicious intention that makes it permissible to kill him (along with the necessity and proportionality of doing so). Some people interpret this as the claim that Villain deserves to be killed, such that killing him in self-defence does not wrong him, or as the claim that Villain’s culpability means that he forfeits any right he has that you do not kill him.

The culpability account can also explain an important feature of self-defence, namely our intuition that should you try to defend yourself against Villain, Villain

cannot then invoke a right of self-defence against *you*. Even though you will now threaten each other, the culpability account will hold that, since you are innocent, you are not a legitimate target of force.

The culpability account is certainly appealing. But it has limitations. For example, consider the following case.³

Elevator

You are in an elevator with Stranger, who is a schizophrenic. Stranger suddenly, and without fault on his part, goes into an extremely violent schizophrenic episode and attacks you. You can stop Stranger from throttling you to death only by killing him.

Again, most people think that you are permitted to defend yourself in a case like *Elevator*. But we cannot explain this permissibility on a culpability account of self-defence. Stranger is not culpable for the attack on you: we can suppose that this is the first time Stranger has had a schizophrenic episode and that he has no reason to think that he might pose this sort of danger to others. Because of his lack of control (or agency) over the threat he poses, Stranger is what is usually described in the literature as an *innocent* or *non-responsible* threatener.

Some accounts of self-defence bite the bullet in cases like *Elevator* and insist that you are not permitted to kill Stranger in self-defence. But the culpability account also rules out using self-defence in a case such as *Vehicle*.⁴

Vehicle

Driver is a very conscientious car-owner, who keeps his car well maintained and who drives very carefully. One day, however, a freak event causes the brakes to fail in Driver's car. Driver loses control of the car and careers towards you as you walk past. He will hit and kill you unless you use your anti-tank gun to blow up both the car and Driver.

Even people who think that you are not permitted to kill a wholly non-responsible attacker like Stranger tend to think that you may defend yourself against Driver. Driver bears at least some responsibility for the threat that he poses to you. But Driver isn't a culpable attacker – he isn't maliciously trying to kill you. Indeed, unlike Stranger, Driver isn't trying to kill you at all. A culpability account will thus have to prohibit your killing Driver in self-defence.

If we think that these limitations on the scope of a right to self-defence are plausible, we might support the culpability account. But if we think that an account of self-defence ought to be able to permit killing people like Stranger and Driver, we will have to look elsewhere for an explanation of when and why self-defence is permissible.

The rights-based account

Rights-based accounts of self-defence begin by assuming that people have a right not to be killed. So if it is permissible to kill a person, we must explain what happens to their right not to be killed, such that killing them does not violate that right. These accounts seek to explain the permissibility of killing a person in self-defence by describing the conditions under which a person can be said to have lost their right not to be killed.

One of the most influential rights-based accounts of self-defence comes from Judith Jarvis Thomson. Thomson argues that the right not to be killed is a claim right. Claim rights confer upon other people both a duty of non-interference and a duty of assistance. A duty is a morally binding obligation to behave in a particular way. If you are under a duty of non-interference with respect to a person's right, you must refrain from interfering with that right. If you are under a duty of assistance with respect to a person's right, you must (help to) ensure that the right is fulfilled.

For example, children have a claim right to an education. This claim is held against their parents (amongst other people). This means that parents have a duty to (a) refrain from preventing their children from going to school and (b) actively ensure that their children go to school. The right not to be killed is a claim right that is held against people in general. It requires both that (a) other people refrain from killing me and (b) other people help me to defend myself if I am unjustly attacked. We will return to this idea of defending others – other-defence, as it is usually called – in more detail later.

Thomson's explanation of how a person loses their right not to be killed turns upon her understanding of what it is to fail in the duty not to kill someone. Thomson explains the permissibility of your using force in *Truck* in the following way. In *Truck*, Villain has no right to kill you. He therefore has a duty *not* to kill you. He will fail in this duty, and violate your right not to be killed, if he does in fact kill you. Because Villain is on course to violate your right not to be killed, he lacks a right that you not kill him. So if you kill Villain, you do not violate his right not to be killed, because he no longer has such a right. This explains why it is permissible for you to kill Villain.

Thomson's view accommodates our intuition that Villain may not fight back against you as you try to defend yourself. In Thomson's view, one may lethally defend oneself against a person only if that person will otherwise violate one's right not to be killed (or violate some other sufficiently important right). Since Villain has lost his right not to be killed, you will not violate Villain's right by killing him. Thus, Villain cannot lay claim to a right to defend himself against you.

Thomson's account also generates the result that it is permissible to kill Stranger in *Elevator*. We might think that, given Stranger's lack of control, it sounds implausible to say that he violates your right by killing you. But on Thomson's account of rights, provided that we agree that Stranger has no right to kill you, which seems plausible, it is possible to say that Stranger will violate your right if he does kill

you. Once we say that Stranger lacks a right to kill you, we are committed, she claims, to saying that Stranger has a duty not to kill you. Stranger fails in this duty just in case he does in fact kill you. And for Thomson, to fail in a duty that one owes to a person is simply to violate their right. On her account, then, one can fail in a duty (and thus violate a right) without being morally responsible for that fact. Thomson's view of self-defence thus denies the need for culpability or even moral responsibility on the part of the person posing the threat. What is crucial for permissible defence is that the person whom you will kill in self-defence will otherwise violate your rights.

It is this criterion of permissible defence that limits the scope of the right of defence on Thomson's account. Consider the following case.⁵

Hungry

Villain has locked you in his dungeon. You are slowly starving to death. One day, Villain lowers a plump baby into the dungeon, along with a knife and fork. You can prevent your death from starvation only by killing and eating the baby. If you do not eat the baby, Villain will remove it, and it will survive.

It seems to me a good test of any account of permissible defence that it should prohibit killing the baby in *Hungry*. Thomson's account passes this test. Even though killing the baby is necessary to save your life, and the harm to the baby is proportionate to the harm that faces you, the baby is not on course to violate your rights. You have a right not to be killed that Villain is violating by starving you to death in his dungeon. But the baby is not violating that right, and thus you may not use force against the baby to avert harm to yourself.

Many people have objected to Thomson's claim that it is possible for a person to violate your right not to be killed even if they are not morally responsible for killing you. Consider the following case.⁶

Ray gun

Falling Person has been blown by the wind down a well, at the bottom of which you are trapped. Falling Person will crush you to death unless you vaporise her with your ray gun. If you do not vaporise her, your body will cushion Falling Person's landing, saving her life.

Falling Person is not even acting as she threatens you. Unlike Stranger, there is no sense at all in which we might say that Falling Person is *trying* to kill you. But it still sounds implausible to say that she has a *right* to kill you. Given this, Thomson's account holds that Falling Person has a duty not to kill you. Just like Villain, she will therefore violate your right not to be killed if she kills you. Falling Person therefore lacks a right not to be killed. You are permitted to lethally vaporise Falling Person on precisely the same grounds as you are permitted to kill Villain and Stranger.

However, if we are cautious of the idea that a schizophrenic can violate a person's rights, we are likely to be downright sceptical of the idea that a person who is helplessly falling can violate a person's rights. Michael Otsuka argues that even if we agree with Thomson that Falling Person will kill you, we cannot think that she violates your rights in doing so:

Surely one would not be tempted to make the same claim about a stone that falls on you after being lifted into the air by a tornado. Could things really be any different if it is an unconscious human being that falls on you after being lifted into the air by a tornado? I do not see how the rights-violating power of such a human object could be any greater than the rights-violating power of a chunk of granite. But talk of rights violations has, I think, gone too far if it is based on a theory that implies that a falling stone can violate a human right.

(M. Otsuka (1994), 'Killing the innocent in self-defense'.
Philosophy and Public Affairs 23, No. 1, 74–94, p. 80)

Falling Person is what is sometimes called an *agency-lacking threatener*. She endangers your life, but she has no control over this fact. Her movements do not proceed from her moral agency in the way that we would normally require if we were to describe her as *acting* rather than simply moving. It is this lack of agency that has led several commentators to compare agency-lacking threats to inanimate objects or animals.⁷ We can agree that falling stones and charging tigers kill people. But do we want to be committed to the idea that stones and tigers violate rights? Otsuka argues that such a view is simply implausible and that we should therefore reject both the permissibility of killing Falling Person and Thomson's account of self-defence.

Otsuka certainly seems to be right that falling stones and charging tigers cannot violate human rights. But many people nonetheless share Thomson's intuition that one may defend oneself against both a violent schizophrenic and an innocent falling person. It seems extremely morally demanding to require that you simply allow yourself to be killed in *Elevator* and *Ray Gun*. Certainly, the law would not require you to refrain from using self-defence in such cases. After all, you are innocent too. Can we defend Thomson against Otsuka's criticism?

Here's a passage from David Rodin, in which Rodin elaborates on precisely what he thinks has gone wrong in Thomson's argument to produce the conclusion that a stone rolling down a hill can violate human rights. Rodin suggests that Thomson is mistaken to infer that simply because something has no right to fall on you, it has a duty *not* to fall on you:

The reason why the stone on the hill has no right to fall on you is that it is just not the kind of entity that could be the subject of rights and duties. From the fact that the stone fails to possess a right to fall in this sense, it obviously cannot be inferred that it has a duty not to fall. The correct conclusion is that the stone has neither a right to fall on you nor a duty to abstain from falling,

because it is not a moral subject at all. But precisely the same is true of the falling fat man. He is a moral subject, but *qua* falling object he is just like the stone, neither the subject of a liberty to fall on you, nor of a duty not to fall. The falling is not something he does so his falling cannot be in violation of any duties he owes you. Hence, in crushing you he violates none of your rights.

(D. Rodin (2002), *War and Self-Defense*, New York: Oxford University Press, p. 86)

I think that most of us would agree with Rodin that the stone cannot be the subject of rights and duties. But if the reason for this is that the stone is not a moral subject at all, this argument fails to show that Falling Person cannot be the subject of rights and duties. After all, if she were really just like the falling stone, it would obviously be true that one is permitted to vaporise her, since vaporising stones is morally unproblematic.

If Rodin wants to resist this implication of his analogy, he must hold that Falling Person is enough of a moral subject to have a right not to be killed. But the portrayal of Falling Person as a sort of quasi-moral subject, with a right not to be killed but no duty not to kill others, might strike us as rather *ad hoc*. Perhaps Thomson's critics cannot have it both ways. If Falling Person is not the sort of thing that can have duties, perhaps she is not the sort of thing that can have a right not to be killed. If so, you will not violate her right if you kill her. If, however, Falling Person *can* have a right not to be killed, there seems to be at least some reason to think that she can have duties as well. But if we grant this, it seems possible that Thomson's argument succeeds, if it is true that Falling Person has no right to kill you, to begin with.

Say that we interpret Thomson's argument as applying only to those things that are moral subjects in general, which seems like quite a plausible restriction of her view. In this interpretation, her argument seems to escape the perils of falling stones and charging tigers. She can avoid the peculiar conclusion that stones have a duty not to fall on people, because stones are not moral agents in general. And Rodin and Otsuka cannot really deny that Falling Person is a moral agent in general if they want to hold onto the claim that Falling Person, unlike the stone, can have rights. This version of Thomson's argument seems to withstand Rodin's and Otsuka's objection.

There might be other reasons, however, to be cautious of Thomson's account of self-defence. For example, even if her account of rights can overcome the sort of objection that Otsuka and Rodin pose, we might wonder whether she is right about what it takes to fail in a duty to someone, and thereby violate their right. Is it enough that Falling Person will kill you to make it the case that she has failed in a duty she owes you? Or must she *intentionally* kill you (or at least kill you through her reckless behaviour) if we are to say that she has failed in this duty?

On the one hand, we usually think that to say that someone ought to do something entails that they are able to do it (summarised in the philosophical slogan of 'ought implies can'). How can it be true, then, that Falling Person has a duty to

refrain from killing Victim if she is not able to so refrain? But on the other hand, it seems perfectly possible to say that, for example, a person forgot to do something that she ought to have done, such as keep a promise. Forgetting, by its very nature, does not proceed from a person's agency (this is why it sounds odd to remonstrate with a person that she *shouldn't* have forgotten). But we don't usually think that simply because a person forgot to keep their promise, they had no duty to keep their promise. Rather, we think that they failed in a duty even though, in one sense, we realise that they could not help doing so. I am not sure which side of the fence we should come down on here, but it is worth bearing in mind that there may be these more general issues surrounding Thomson's conception of rights, even if her account of self-defence can overcome the specific objection posed by Otsuka and Rodin.

We also noted earlier that Thomson can explain why Villain in *Truck* cannot employ counter-defence against you, because on her account Villain lacks the right not to be killed. But since Thomson's account of permissible defence does not distinguish between culpable attackers and innocent attackers, she cannot allow Falling Person to fight back either. Even though Falling Person is morally innocent, Thomson's account will not allow her to defend herself against you. This implication of Thomson's view might push our intuitions too far if we feel that Falling Person's innocence ought to make at least some difference to her moral status. Even if we think that you may defend yourself against Falling Person, we might think that her innocence at least grants Falling Person the right to fight back. But it isn't clear how we can allow for such a permission if we endorse Thomson's account.

The responsibility account

The *responsibility account* of self-defence is an alternative to both the culpability account and Thomson's rights-based account. Jeff McMahan is the most influential contemporary proponent of this account. McMahan understands permissible defence as foremost a matter of justice. An account of self-defence should tell us who, as a matter of justice, ought to bear a harm, given that someone has to bear it. For example, consider *Truck* again, where Villain is intent on running you down. In this case, either you must kill Villain or Villain will kill you. The question, then, is how this harm of death ought to be distributed between you and Villain. It seems obvious to most of us that it is only fair that Villain be the one to bear the harm. He is, after all, culpably attacking you. McMahan agrees, arguing that justice requires that Villain bear the harm, and thus you may kill him rather than bear the harm yourself.

However, in McMahan's view, it is not, strictly speaking, Villain's culpability for the threat that makes it fair that he bear the harm. Rather, what matters is Villain's moral responsibility for the threat. Of course, Villain is also culpable for the threat. But whilst culpability entails moral responsibility, the converse is not true. Recall the *Vehicle* case, in which Driver helplessly loses control of his car.

Even though Driver is not culpable for posing a threat, McMahan argues that he is morally responsible for doing so.⁸ Driver chose, voluntarily, to engage in driving, knowing that driving poses a small but genuine risk of serious harm to others. McMahan argues that Driver's moral responsibility for the fact that someone must now bear a harm makes it permissible for you to ensure that Driver suffers the harm rather than you.

McMahan summarises this claim by saying that Driver is liable to defensive killing. To say that a person is liable to defensive killing is to say that killing them does not wrong them and does not violate their rights. A person who is liable to be killed has no justified complaint against being killed. McMahan thinks that one can become liable to defensive killing only if one is morally responsible for an unjust threat. An unjust threat is a threat of harm that will wrong the person upon whom it is inflicted: it is a threat of harm directed at a person who is not liable to that harm. So Driver poses an unjust threat to you in *Vehicle* because you are not unjustly threatening anyone, and are thus not liable to be killed. If he kills you, Driver will inflict an unjust harm upon you. And because he is morally responsible for posing this unjust threat, he will wrong you in killing you. It is thus Driver's moral responsibility for an unjust threat that grounds your permission to kill Driver in self-defence. The same is true of Villain in *Truck*.

However, notice that this account will not generate a permission of self-defence in cases such as *Elevator* or *Ray Gun*. It is true that you are not liable to be killed in either of these cases. Stranger and Falling Person thus pose unjust threats to you. But neither Stranger nor Falling Person is morally responsible for the unjust threat that they pose, and therefore they are not liable to be killed by you in self-defence. In McMahan's view, neither Stranger nor Falling Person will wrong you if they kill you. But you will wrong *them* if you kill them, because they are not liable to be killed. Hence, the defensive threat that you pose will be unjust, and you will be morally responsible for inflicting it because your defence will be deliberate. So whilst the responsibility account is more permissive than the culpability account, it is more restrictive than Thomson's rights-based account.

The doctrine of double effect

The final account of self-defence that we will consider is based upon the doctrine of double effect (DDE). The DDE is not exclusively an account of permissible self-defence: double-effect reasoning is used extensively in moral theory, including just war theory. The DDE finds its roots in Catholic theology and is an attempt to reconcile the Catholic ideal of 'doing no harm' with our intuition that inflicting defensive harm is sometimes permissible. Of course, one need not hold with Catholicism to think the DDE plausible: secular writers have made frequent use of double-effect reasoning to explain why it can sometimes be permissible to cause harm.

The DDE rests upon an alleged distinction between what I *intend* to cause by my actions (or what I intend to allow through my inaction) and what I merely *foresee* I will cause by my actions (or what I foresee I will allow through my

inaction). There are, undeniably, some cases in which it seems plausible to distinguish between effects I intend and effects I merely foresee. For example, the more electricity I use, the larger the electricity bill I run up. But it would be odd to insist that I turned on the light with the intention of running up a large electricity bill. Rather, I intend to improve the lighting in the room, and merely foresee that I will receive a larger bill as a result. Here, the lighting of the room and the increasing of my bill seem to be separable effects of the action of switching on the light, such that I can properly be said to intend one but not the other. Proponents of the DDE often emphasise this separateness by pointing out that in the lucky event that I receive a small electricity bill, I will not be disappointed or feel that I have been unsuccessful in achieving some goal of mine. The DDE can explain this, because on this view I had no intention of running up a large bill, and thus my failure to do so is not a frustration of my intentions.

Of course, the DDE does not hold only that we can conceptually distinguish the effects that I intend from the effects that I foresee. It also holds that there is a morally significant difference between these two effects. The DDE claims that whilst it is always impermissible to intend harm, it is nonetheless sometimes permissible to cause harm as a foreseen side-effect of an action, if doing so is necessary to achieve some proportionate good. The double-effect explanation of self-defence, then, insists that it would be wrong to intend harm to one's attacker. However, causing harm to the attacker can nonetheless be permissible provided that it is merely a foreseen side-effect of pursuing the good end of saving one's life. In other words, the permissibility of self-defence lies in the fact that the defender does not aim at doing harm, but aims only at the morally permissible end of saving their life.

Double-effect reasoning has been similarly employed to argue for the permissibility of other sorts of killing, such as euthanasia. According to the DDE, it would be morally impermissible for a physician to intend to kill their patient by giving them a lethal dose of some drug. But it is permissible to intend to relieve the patient's suffering by giving them a lethal dose of the drug, even if the physician foresees that the patient will be killed as a side-effect.

Unsurprisingly, this divorcing of intended consequences from foreseen consequences has attracted a great deal of criticism. In some cases, the two effects of an action – the good end and the harmful side-effect – seem too closely entwined for one to plausibly claim that one intended to produce only the good effect. Consider the *Ray Gun* case again. The 'closeness' objection insists that it is wholly implausible that you might vaporise Falling Person, but intend only to save your life, and not to kill her. Sustaining the intended/foreseen distinction is particularly difficult in such cases because the harm is what achieves the good. That Falling Person is vaporised is what makes it the case that your life is saved. There are not two things here: the saving of your life and the vaporising of Falling Person. Rather, your being saved consists in her being vaporised, and to vaporise her is to kill her. It is quite a stretch of our ordinary language to describe the killing of Falling Person as a side-effect of an action that saves your life.

It can also be difficult, particularly in cases of self-defence, to meet the DDE's requirement that the harm not be intended as a means. Imagine that the only way that I can stop you from attacking me is to cause you so much pain that you have no choice but to desist. When I stab you in the leg, causing you pain seems to be clearly instrumental in saving my life. It is not a mere side-effect of stabbing you: it is the reason why I stab you, because your being in pain is the means to my escaping. Stabbing you seems clearly permissible even though I intend to cause you pain. But the DDE holds that harm is permissible only if it is unintended: the harm must not be a means of achieving the end.

So there are what look like instances of permissible defence that the DDE cannot allow. And there are what look like instances of impermissible defence that the DDE seems to permit. Consider the following case.⁹

Bomb

A runaway trolley is heading towards you as you lie trapped on the trolley tracks. You will be killed unless you throw a grenade at the trolley and destroy it. However, the blast from the grenade will also kill an innocent person who is lying unconscious by the side of the tracks.

Most accounts of self-defence hold that it is impermissible for you to throw the bomb if doing so will kill the innocent person. But the innocent person's death would not be intended. It would be merely a foreseen side-effect of saving your life. How can proponents of the DDE prohibit killing this innocent person?

They might do so by arguing that this killing fails the proportionality requirement. How so? Well, the death of one innocent person certainly seems comparable to the death of another innocent person, so they could not do so on the threshold interpretation of proportionality that I suggested earlier. But they might opt instead for the simple outweighing model of proportionality, and insist that one may inflict harm only if the good achieved is in excess of the harm caused. Since in *Bomb* the harm inflicted would be equal to the harm avoided, the killing would be prohibited by this understanding of proportionality. This strategy would prohibit killing innocently threatening people, such as Falling Person. But we might think this a reasonable stance: perhaps such killings *are* impermissible.

Of course, if proponents of the DDE adopt this strategy, they will have to endorse the idea that the life of an innocent person is worth more than the life of a culpable person. Only in this way could the DDE permit the killing of a culpable attacker to save one's own life. But, again, we might think that this is correct. Perhaps harming a culpable person is not as morally bad as harming an innocent person because harm to the culpable person counts for less.

However, this explanation of the impermissibility of throwing the grenade in *Bomb* might still be problematic, because it is now unclear whether the DDE can permit you to kill a culpable attacker to prevent a lesser harm than death. Recall my suggestion earlier that lethal force seems like a proportionate response to a threat of paralysis or kidnapping. If it is indeed permissible to kill a culpable attacker to

prevent such harms, the doctrine of double effect must hold that these harms are worse than the death of a culpable person. We now seem to be attributing significantly less weight to the culpable person's life on this account. Furthermore, most people think that I may kill more than one culpable attacker to prevent them from inflicting serious harm upon me. But, if this is true, this means that the life of a culpable person must be worth very little indeed, if I may kill numerous would-be kidnappers rather than let myself be kidnapped.

Mark Timmons suggests that, at root, the DDE does nothing more than offer alternative, but equally true, descriptions of the same action. The permissibility of the action then depends on which description one chooses to endorse, as in the euthanasia case described earlier:

So describing the action of the physician in one way yields the verdict that her action is an intentional killing, but when the action is described in another way, the [DDE] yields the verdict that her action is not an intentional killing. This means that the verdict we get using the [DDE] depends on how we describe the physician's action. Clearly something is wrong here. Either the physician's action is morally wrong or it is not morally wrong. It is nonsense to say that her action is wrong if described one way but not wrong if described another way, where both descriptions apply truly to her action.

(M. Timmons (2002), *Moral Theory: An Introduction*, Lanham, MD: Rowman and Littlefield, p. 91)

Timmons' complaint is echoed by many writers. What ought to concern us is the moral permissibility of an action that will inflict a certain harm. But the DDE seems to tell us that whether an action is permissible depends upon how we describe it. On some descriptions, the action will be permissible. But on other, no less accurate, descriptions the action will be wrong.

There lurks behind this complaint a deeper suspicion that although the DDE might be used to justify a given action, it does not really explain *why* the action is permissible. Rather, it simply offers support for our prior convictions about an action's permissibility. When the end and the 'side-effect' are closely entwined, we can adopt the description that best suits our purposes. If we think that killing in self-defence is permissible, then we can adopt the description that pronounces the harm as an unintended side-effect. If we think it is impermissible, then we can insist that the harm is intended. The DDE offers us no independent argument for preferring one description over the other – that is, for judging whether something really is a side-effect or an end – but it is this that would tell us whether or not killing in self-defence is morally permissible. More generally, we might feel that the DDE cannot do justice to the idea that harm caused in self-defence is morally permissible even when it is intended. The DDE fails to capture the sense that I am allowed to kill a person who will otherwise kill me, and that I need not try to justify or excuse my killing by invoking some different goal or re-description of my action that the DDE will deem permissible.

As we can see, attempts to use double-effect reasoning to explain self-defence are beset with problems. However, in other areas of moral theory, the DDE is perhaps more persuasive. Certainly, it seems sometimes to make a moral difference whether I intend to cause something, or merely foresee that I will cause it. The famous trolley problem,¹⁰ for example, is often explained by invoking the DDE.

Trolley

A runaway trolley is heading towards five people who cannot get out of the way before it hits them. The trolley will kill them if it hits them. You are standing next to a switch that you can use to divert the trolley down a side track. However, the trolley will then kill a single individual who is stuck on the side track.

Most people think that you may divert the trolley. An obvious explanation of the permissibility of doing so is the difference in the numbers involved. Five people will die if you do nothing, but only one person will die if you divert the trolley. But if numbers are doing all the work, we commit ourselves to all sorts of unwelcome conclusions.

For example, consider *Transplant*.¹¹

Transplant

The trolley has hit the five people, and they are all badly injured. In order to survive, each person needs a different organ transplant. Luckily, the one on the side track is unharmed. You, the surgeon, can cut up the one and distribute his organs amongst the five, saving their lives. The one will, of course, die.

In this case, very few people think that you may save the five people at the cost of the one's life. But if what makes diverting the trolley permissible is simply the difference in numbers, there should be no difference between the two cases. This tells us that it is not merely the numbers doing the moral work in the original *Trolley* case. What might be doing the moral work, however, is the DDE. In *Trolley*, you do not intend to kill the one. You would, after all, divert the trolley even in his absence. And, as Philippa Foot says, should the one miraculously survive being hit by the trolley, you would not pick up a spare bit of track and proceed to beat him to death with it.¹² Killing the one is not part of your aims: all that you aim at is the (permissible) end of saving the five.

It is hard to make this story stick in *Transplant*, though. Here your plan to save the five depends upon the presence of the one. Without him, there is no rescue plan. And so you cannot very well claim to have merely foreseen his death, since killing him is part of the plan. Of course, we might argue that you don't intend to kill him. All you intend to do is remove his organs. You merely foresee, with an appropriate sense of regret, that he will die if you do this. But now we run into

the closeness objection again. The closer the end to the alleged side-effect, the harder it becomes both to separate those effects and to plausibly postulate some morally significant difference between that which is intended and that which is merely foreseen. Removing the one's organs constitutes killing him, and thus it's hard to see how the intended/foreseen distinction can have any purchase in this case.

So what one intends seems relevant when dealing with some moral problems, but not others. With respect to just war theory, the DDE is primarily used to justify the collateral killing of non-combatants. Whilst non-combatants are not usually regarded as legitimate targets in war, just war theorists have often argued that killing non-combatants as a foreseen side-effect of attacking a military target can be morally permissible. We will consider this application of the DDE in Chapter 7.

Other-defence

An important aspect of any theory of self-defence is what it tells us about the permissibility of other-defence. Other-defence occurs when a third party comes to the aid of a person who is under attack. We usually think that if a person is unjustly attacked, it is morally permissible for others to come to their rescue. We might even think that rendering such assistance is, at least *prima facie*, morally required. We generally condemn those who stand by and do nothing when a person is under attack, which suggests that we conceive of other-defence as a duty rather than a mere permission. Indeed, the early Christian just war theorist St Ambrose believed that defensive force could be permissibly employed *only* in defence of other people's lives and that such force could be 'actually demanded on moral grounds'.¹³

Given this, our account of other-defence has the potential to play an important role not only in individual defence but also in the context of war. There are two readily apparent ways in which other-defence is relevant to war. One is in terms of combatants who are engaged in a defensive war against aggression. The combatants might well be understood to be defending themselves. But they are also defending the other citizens of their nation. For writers like Ambrose, this is crucial to the permissibility of the war. It also has significant implications for things such as proportionality in war. If the threat that I am averting is a threat to many lives, the amount of force that I may employ in averting that threat will be correspondingly greater.

The second obvious role that other-defence can play in the ethics of war is in explaining the permissibility of humanitarian wars of intervention. There is a growing consensus amongst just war theorists, and people in general, that it is not only when a nation wages a defensive war that its use of military force is permissible. We also think it can be permissible for a nation to use its armed forces to defend citizens in another country against either external or internal threats. If

the permissibility of such wars is explained by the fact that they are instances of other-defence, we can perhaps explain, for example, the condemnation that followed the failure of countries like the United States to intervene in Rwanda in 1994. Since other-defence is often conceived of as a duty, we have good grounds for thinking that wars of intervention are not only morally permissible but also morally required.

Consider the following case.

Rescue

Villain is chasing Damsel because he hates her and he wants to kill her. Poor Damsel has no means of defence. Luckily, Knight comes riding past. Knight can easily kill Villain and save Damsel's life.

I expect that nearly all of us will think that Knight is morally permitted to save Damsel by killing Villain. And many of us might think that not only is Knight permitted to do this, but he is under obligation to do this. He would act wrongly if he simply kept riding and ignored Damsel's plight when he could so easily rescue her.

The obligatory nature of other-defence is often explained by the idea that Damsel has a right not to be killed. As I described earlier, people usually understand the right not to be killed as a claim right that confers duties of non-interference and of assistance on other people. Damsel is permitted to defend herself against Villain because, in trying to kill Damsel, Villain fails in the duty of non-interference that he owes her. Doing so violates Damsel's right. And if Knight fails to assist Damsel, he fails in the duty of assistance that he owes her, and also violates her rights.

But what should we say about *Retreat*?

Retreat

Villain is chasing Damsel because he hates her and he wants to kill her. Poor Damsel has no means of defence. Knight comes riding past. Knight knows Villain, and knows what a skilled swordsman Villain is. Knight knows that his chances of defeating Villain without suffering serious injury to himself are very slim. Knight can safely retreat without attracting Villain's attention.

Is Knight required to try to save Damsel even at serious risk to himself? Most people think that he is not so required. This suggests that even if other-defence is a duty, rather than a mere permission, it is a defeasible duty. Our duty to assist is not so stringent as to require us to bear a significant risk of serious harm to help others (although it might require that we bear some risks of harm).

We might invoke the idea of an agent-relative reason to explain why this is the case. An agent-relative reason is a reason that has genuine weight for the agent, but that does not have weight for others. For example, imagine that a man can either

save his wife or a stranger from drowning. The fact that one of the drowning people is his wife gives the man a very good reason to save her instead of the other person – we might even think that it gives him a morally conclusive reason. But he could hardly expect this same fact – that she is *his* wife – to give the coastguard a good reason to save her rather than the other drowning person. From the coastguard's perspective, this agent-relative reason is no reason at all.

Similarly, we might think that when it comes to a choice between saving his own life and saving Damsel's life, Knight has good reason to choose to save his own life. His agent-relative reason to prefer his own survival justifies his decision to fail to intervene on Damsel's behalf. The reason is not so strong that it could justify his deciding to let Damsel be killed rather than suffer some minor harm in rescuing her. But, at some point, the risk to Knight will be sufficiently great that he may permissibly refuse to intervene rather than take that risk. He would still be permitted to intervene, of course. But such intervention would be an act of supererogation, not of obligation.

We can also think of cases in which a third party might not have even a permission of other-defence, let alone a duty. Recall the *Ray Gun* example, in which you are about to be crushed to death by Falling Person. In the discussion of Thomson's view, I drew attention to the fact that Thomson denies that Falling Person may defend herself against you. Falling Person lacks a right not to be killed that could ground her right to self-defence. And on the very same grounds, Thomson affirms a permission for third parties to kill Falling Person on your behalf. Since you, unlike Falling Person, are about to suffer a violation of your right not to be killed, a third party may intervene to prevent this violation.¹⁴

Whether we think that this is correct will probably depend on whether we agree with Thomson about the impermissibility of Falling Person's defending herself against you. If we share Thomson's view that Falling Person is not permitted to use counter-defence against you, we might be amenable to the idea that third parties are permitted to kill Falling Person on your behalf. But if we are hesitant about this first claim, I suspect we will be similarly hesitant about what third parties may do to Falling Person. We might think that part of what explains your permission to kill Falling Person is the existence of an agent-relative reason to prefer your own survival, just as we might think that Knight is allowed to prefer his survival to Damsel's. But since this is an agent-relative reason, it can have no weight for a third party. If this is right, then third parties lack sufficient reason to choose your life over that of another, equally innocent, person. But, again, this seems to undermine Thomson's view that it is rights-violation that does all the work in self-defence.

Chapter summary

Most people think that it is permissible to use defensive force against an attacker, provided that the force is necessary and proportionate. But these two restrictions on the use of force are only part of the story. Rival accounts of self-defence aim

to give us a more complete understanding of self-defence. Culpability accounts invoke the fault of an attacker to explain why defensive force is permissible. Rights-based accounts rely on showing that an attacker has lost the right not to be killed. Responsibility-based accounts like McMahan's focus on liability to defensive killing, where this is acquired through moral responsibility for an unjust threat. And accounts based upon the doctrine of double effect argue that defensive force is permissible if, and only if, the defender merely foresees the harm that she causes. As we have seen, these different accounts generate different scopes of permissible defence.

While self-defence is usually conceived of as a right or permission, other-defence is usually thought to be a duty. But this duty is fairly limited, on most accounts. If I can easily rescue someone from unjust harm, I usually have a duty to do so. But I do not have to expose myself to a significant risk of serious harm in order to save someone else from an even greater harm. And there may be cases in which a person is permitted to defend themselves but a third party is not permitted to assist them. Thus, while the duty of other-defence is seemingly grounded in the victim's own right to self-defence, it does not follow that I may do in other-defence whatever they may do in self-defence.

Questions for discussion

- 1 Is it permissible to kill Stranger in *Elevator*? Why?
- 2 Which account of self-defence strikes you as the most plausible?
- 3 Do you agree that other-defence is a duty, or is coming to the aid of others always supererogatory?
- 4 How plausible is the doctrine of double effect?
- 5 Is it permissible to throw the grenade in *Bomb*? Why?

Notes

- 1 Some parts of the USA have an even more permissive stance, having passed so-called 'Stand Your Ground' laws that do not require retreat even in public places.
- 2 Based on Judith Jarvis Thomson's case in 'Self-defence' (1991), *Philosophy and Public Affairs* 20, No. 4, 283–310, p. 283.
- 3 Based on George Fletcher's case in 'Proportionality and the psychotic aggressor' (1973), *Israel Law Review* 8, 367–90, p. 371.
- 4 Based on Jeff McMahan's case in 'The basis of moral liability to defensive killing' (2005), *Philosophical Issues* 15, Normativity, 386–405, p. 393.
- 5 Based on Thomson's case in 'Self-defence', p. 291.
- 6 Based on Robert Nozick's case in *Anarchy, State and Utopia* (1974). Malden, MA: Basic Books, p. 34.
- 7 See M. Otsuka (1994), 'Killing the innocent in self-defence', *Philosophy and Public Affairs* 23, No. 1, 74–94; J. McMahan (2002), *The Ethics of Killing: Problems at the Margins of Life*. New York: Oxford University Press, p. 409; and D. Rodin (2002), *War and Self-Defense*. New York: Oxford University Press, pp. 81–3.
- 8 McMahan, 'The basis of moral liability to defensive killing', p. 394.
- 9 Otsuka, 'Killing the innocent in self-defence', p. 85.

- 10 P. Foot (1967), 'The problem of abortion and the doctrine of double effect', *Oxford Review* 5, 5–15.
- 11 J. J. Thomson (1986), 'Self-defense and rights', in VV. Parent (ed.), *Rights, Restitution and Risk*. Cambridge, MA: Harvard University Press, p. 34.
- 12 Foot, 'Abortion and double effect', p. 9.
- 13 L. J. Swift (1970), 'St. Ambrose on violence and war', *Transactions and Proceedings of the American Philological Association* 101, 533–43, p. 535.
- 14 Thomson, 'Self-defense', p. 308.

Suggested reading

Judith Thomson (1991), 'Self-defense', *Philosophy and Public Affairs* 20, No. 4, pp. 283–310.

One of the classic papers in the self-defence literature, this paper defends a rights-based account of self-defence that permits the killing of innocent threateners.

Michael Otsuka (1994), 'Killing the innocent in self-defence', *Philosophy and Public Affairs* 23, No. 1, pp. 74–9.

This paper offers the most influential critique of Thomson's account, arguing against a permission to kill innocent threateners.

Noam Zohar (2004), 'Innocence and complex threats: Upholding the war ethic and the condemnation of terrorism', *Ethics* 114, pp. 734–51.

This paper looks specifically at the implications of Thomson's account of self-defence for the ethics of war.

Jeff McMahan (2005), 'The basis of moral liability to defensive killing', *Philosophical Issues* 15, Normativity, pp. 386–405.

This paper outlines McMahan's criticisms of several accounts of self-defence, and develops his own responsibility-based account. Given the importance of McMahan's work in just war theory, it's a good idea to read this article.

The Ethics of War and Peace

An Introduction

Third Edition

Helen Frowe

Cover image: guvendemir/Getty Images

Third edition published 2023

by Routledge

4 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

and by Routledge

605 Third Avenue, New York, NY 10158

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2023 Helen Frowe

The right of Helen Frowe to be identified as author of this work has been asserted by her in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

First edition published by Routledge 2011

Second edition published by Routledge 2015

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data

A catalog record for this book has been requested

ISBN: 978-1-032-23055-9 (hbk)

ISBN: 978-1-032-23054-2 (pbk)

ISBN: 978-1-003-27546-6 (ebk)

DOI: 10.4324/9781003275466

Typeset in Times New Roman

by Apex CoVantage, LLC