

UEHIRO SERIES IN PRACTICAL ETHICS

General Editor: Julian Savulescu, University of Oxford

Choosing Children

The Ethical Dilemmas of Genetic Intervention

Jonathan Glover

Messy Morality

The Challenge of Politics

C. A. J. Coady

Killing in War

Jeff McMahan

Killing in War

JEFF MCMAHAN

CLARENDON PRESS · OXFORD

2009

Assume, then, that moral responsibility for an unjustified threat is the basis of liability to attack in war.² There are various forms and degrees of responsibility, and therefore also of liability. In particular, the degree of a person's responsibility for unjustifiably posing a threat of wrongful harm to others varies with the significance of any excusing conditions that may apply to his action. The more a person is excused for some objectively wrongful act, the less responsible he is for the consequences, and the less liable he may be to defensive action to prevent those consequences from occurring. If this is right, it could be highly relevant to the morality of killing in war. If, for example, an unjust combatant is fully excused for fighting in an unjust war, that may mean that his liability to defensive action is comparatively weak. In that case there might be a requirement that just combatants exercise certain forms of restraint in fighting against him. Or if there were no such requirement, the justification for attacking him without restraint might have to appeal at least in part to considerations other than his liability to attack, and those other considerations would have to be identified.

To try to elucidate the relation between excuses and liability, I will focus initially on cases of individual self-defense outside the context of war. Some of the examples I will use for purposes of illustration will presuppose that agents have knowledge that it is sometimes difficult to have even in cases of individual self-defense and that it is virtually impossible to have in the complex and confused circumstances of war. But it is helpful to try to get clear about what is at issue in simplified cases, even if they have certain idealized features, in order to know what to look for in the more complicated cases involving combatants in war that are our ultimate concern here. The conclusions we may draw from simplified cases of individual self-defense may have no direct application in war because the conditions for knowledge are so different. But it is important to identify the considerations that would be relevant in war if only we could have knowledge of them, in order to determine how we might most effectively try to compensate for the absence of that relevant knowledge.

It will help to have before us a set of distinctions among various different kinds of agent, to whom I will refer as "Threats," who pose a threat to others. I will define certain categories of Threat, offering one or more hypothetical examples of each. This will, I hope, introduce some conceptual clarity, which will then enable us to explore questions about whether, why, and to what extent different types of Threat are liable to defensive action, after which we can try to determine into which

category or categories most unjust combatants fit, with the ultimate aim of better understanding their moral status in war.

4.1.2 Culpable Threats

I begin with the category at one end of the spectrum: *Culpable Threats*. These are people who pose a threat of wrongful harm to others and have neither justification, permission, nor excuse. They may intend the harm they threaten, or the risk they impose, or the threat may arise from action that is reckless or negligent. But because they have neither justification nor excuse, they are fully culpable for their threatening action. As such, they are fully liable to necessary and proportionate defensive action. A man who attempts to murder his wife so that he can inherit her money is a Culpable Threat. Not only the wife herself but any third party would be fully justified in killing this man if that were necessary to prevent him from killing her—or even just the surest means of saving her.

It is generally agreed that the proportionality restriction on killing a Culpable Threat is weaker than it is in other cases. Precisely because of the Culpable Threat's full culpability, it can be proportionate to inflict a significantly greater harm on a Culpable Threat if that is necessary to prevent him from inflicting a lesser harm on an innocent victim. The proportionality restriction is thought to be particularly weak in the case of lethal threats. It is, for example, often claimed that it can be permissible to kill any number of Culpable Threats when that is necessary to prevent them from killing a single innocent person. Others, though not as many, think that if a Culpable Threat intends to kill an innocent victim, it can be permissible to kill him if that is necessary to avert the threat no matter how low the probability of his succeeding in killing his victim would otherwise have been.

4.1.3 Partially Excused Threats

Culpable Threats are *fully* culpable; there are no excusing conditions that apply to their action. It is possible, however, to be culpable for an act while having a partial excuse. This is obvious, for when an excuse is only partial, some residue of culpability remains. There is therefore a large category of what I will call *Partially Excused Threats*: people who unjustifiably pose a threat of wrongful harm to others but whose action is excused to some extent, though not fully. A Partially

Excused Threat may have an excuse of any of the three broad types I identified earlier. A person might, for example, unjustifiably pose a threat to an innocent person only because he has himself been credibly threatened with some sanction if he fails to harm that person. If the level of duress to which he is exposed could be resisted by a person of ordinary fortitude, and in particular if it is insufficient to overwhelm his own will, his excuse is only partial, only mitigating. The strength of the excuse is a function of, among other things, the severity of the sanction and the magnitude of the harm he threatens to cause.

A person might also unjustifiably pose a threat to an innocent person by mistake, because he lacks relevant knowledge. But if his ignorance was avoidable—if, for example, he has been negligent in investigating the facts relevant to the permissibility of his action—his excuse is only partial. The same is true if he acts with a diminished capacity for responsible agency, but not in the complete absence of that capacity—as, for example, in the case of the person who threatens to destroy your valued possession. The drug he was involuntarily administered has impaired his capacity for self-control, but not eliminated it.

Partially Excused Threats are not necessarily less culpable than Culpable Threats. A Culpable Threat may, for example, be fully culpable for intentionally imposing a comparatively minor threat, while a Partially Excused Threat may be excused to some degree for negligently posing a greater threat. If the threat posed by the Partially Excused Threat is significantly greater, he may be more culpable than the Culpable Threat, despite the fact that his offense is negligent rather than intentional, and that he has a partial excuse. For the degree of an agent's culpability is a function of all these variables—whether the wrongful threat is intentional, reckless, or negligent, whether the agent has an excuse and how strong that excuse is, and the magnitude of the threatened harm—as well as others.

The proportionality constraint on defensive action against a Partially Excused Threat is more stringent than it is in the case of a Culpable Threat. This does not mean, of course, that it is always permissible to inflict greater harm on a Culpable Threat. What it means is that the extent of a Partially Excused Threat's liability is always discounted for his excuse—that is, that his liability would be greater without the excuse, if all other considerations remained the same. In other words, his liability would be greater if he were a Culpable Threat, and other things were equal.

As I noted earlier, the difference in liability emerges in the narrow proportionality constraint. Suppose that someone is attacking you and that if you do not defend yourself, you will be severely injured, though not killed. You have two defensive options. You can kill the Threat, thereby escaping entirely unharmed. Or you can incapacitate him in a way that will injure but not kill him—but only after he will have inflicted a lesser though still substantial injury on you. Suppose that if the person attacking you is a Culpable Threat, it would be proportionate to kill him. In that case, it may be permissible to kill him *rather than* incapacitate him at the cost of suffering a substantial injury—for example, a broken arm. Yet if you know that he has a partial excuse, it may be wrong for you to kill him—that is, you may be morally required to suffer the broken arm in order to avoid killing him. Suppose, for example, that he mistakenly believes that it is his duty to kill you. It is, in the circumstances, a natural mistake, but he could have avoided it if he had taken greater care in investigating the facts before he acted. If the excuse significantly diminishes his culpability, killing him may be disproportionate. In short, the harmfulness of the defensive action to which the Partially Excused Threat is liable varies with the degree of his culpability.

One may wonder why a wholly innocent victim might be required to share the cost of a Partially Excused Threat's wrongful action. Why should not the Threat be required to suffer the entire cost? Or suppose that you could also save yourself without killing the Threat, by breaking some other innocent person's arm. Why should *you*, among all the innocent people in the world, be singled out as the one who has to share the cost of his wrongful action with the Partially Excused Threat?

The answer, I think, is that for you the alternative to sharing the cost with the Threat is to *kill* him, and to do so intentionally, and the option of dividing the cost between the Threat and another innocent person requires you intentionally to break an innocent person's arm. In short, the explanation of why *you* have to share the cost appeals to the distinction between doing and allowing. To see this, suppose that a stranger is drowning and the only way you can save him would involve breaking your arm. Our general practice of refusing to make small sacrifices, such as sending money to Oxfam, to save people who will otherwise die, suggests that we do not believe that it is obligatory to *save* a stranger at the cost of suffering a broken arm. So the reason why you might be required to accept a broken arm for the sake of the Partially Excused Threat, who is not even innocent in the way the

drowning stranger is, must be that what you must do to avoid it is to *kill* someone rather than merely to allow someone to die.

4.1.4 Excused Threats and Innocent Threats

In the next category are those who unjustifiably pose a threat of wrongful harm to others but are *fully* excused for doing so. The paradigm case of an *Excused Threat* is a person who acts under irresistible duress—that is, the sanction he faces, or the harm he will suffer, if he does not unjustifiably threaten someone else is so severe that it overwhelms his will, and would overwhelm the will of anyone else with a normal capacity for the exercise of willpower.

It is important to note that the claim that a person is fully excused for an act of objective wrongdoing implies only that the person is not culpable, that he or she is entirely blameless. It does *not* necessarily imply that the person is absolved of all *responsibility*. A person may be responsible for his objectively wrongful action even if he is not blameworthy. This is true, for example, in most if not all cases of what we call irresistible duress. When we say that duress is irresistible, we usually do not mean that literally. We concede that *some* people could and indeed would resist, and that it was physically and in some sense psychologically possible for the person who failed to resist to have resisted instead. There is therefore a basis for holding him responsible. But the standard of responsibility is not the standard of culpability. We do not accept that all those who are responsible for acting wrongly are also blamable. There are, of course, *some* cases in which what counts as a full excuse on an objective account of permissibility absolves a person not only of all culpability but also of all responsibility. But not all cases of full excuse are like this. I will say more about this later.

I have thus far focused on duress as an example of a full excuse. This was deliberate, for what I have called epistemically-based excuses pose a problem for the taxonomy of Threats. Consider a person who poses a lethal threat to another on the basis of factual and moral beliefs that he is fully epistemically justified in having—that is, he is justified in having these beliefs and in assigning them a degree of credence approaching certainty. But the beliefs are in fact false. If they were true, he would be objectively justified in killing the person he is now attempting to kill, who is in fact wholly innocent. On an objective account of justification, he is acting impermissibly but has a full epistemic excuse. But on a subjective account of justification, which makes a person's

justified beliefs relevant to the permissibility of his action, this person acts permissibly and indeed justifiably, and therefore cannot have an excuse, for there is nothing for him to be excused *for*.

Call this person an *Epistemically Justified Mistaken Threat*. According to an objective account of permissibility, or an objective account of justification, he is an Excused Threat. But not according to a subjective account. It seems, therefore, that we need another label for such a Threat for those who accept subjective accounts of permissibility and justification. Let us say, then, that on a subjective account, he is an *Innocent Threat*. An Innocent Threat is someone who in objective terms acts impermissibly in posing a threat to another but also acts subjectively permissibly, or even with subjective justification. (This is in two respects an unfortunate label. First, the term "Innocent Threat" has several meanings in the literature, all different from the one I propose. Second, I have until now followed the just war tradition in using "innocent" to mean, roughly, "not liable." But it is possible—indeed I will argue that it is true—that some Innocent Threats in my sense *are* morally liable to defensive attack. But because the other labels I have considered are at least equally problematic, I will persist in using the term "Innocent Threat" in this particular technical sense.)

The categories of Excused Threat and Innocent Threat are therefore overlapping but not coextensive. According to both objective and subjective accounts of permissibility, a Threat who poses an objectively unjustified threat on the basis of irresistible duress is only an Excused Threat. But a person who poses an objectively unjustified threat on the basis of what the Scholastics called invincible ignorance is an Excused Threat according to an objective account but an Innocent Threat according to a subjective account.

When a Threat is excused on one account but acts permissibly on the other, the different labels refer to the same facts. In such cases of overlap, there can obviously be no substantive moral difference between an Excused Threat and an Innocent Threat. There are, however, relevant differences within the category of Innocent Threats (and, of course, corresponding differences within the category of Excused Threats, on an objective account). One such difference is that between those who act merely permissibly, with no positive moral reason, and those who act for a positive moral reason, and thus act with subjective justification. There are also relevant differences among those who intend to cause harm, those who knowingly but unintentionally cause harm, those who foreseeably risk causing harm, and those who cannot foresee that their

action will cause harm. All such Threats may be acting in ways that are objectively impermissible but subjectively permissible.

It will be helpful to have particular examples that exemplify some of the various possibilities. Here are some cases that illustrate the distinctions that can be drawn.

(1) *The Resident* The identical twin of a notorious mass murderer is driving in the middle of a stormy night in a remote area when his car breaks down. He is nonculpably unaware that his twin brother, the murderer, has within the past few hours escaped from prison in just this area, and that the residents have been warned of the escape. The murderer's notoriety derives from his invariable *modus operandi*: he violently breaks into people's homes and kills them instantly. As the twin whose car has broken down approaches a house to request to use the telephone, the resident of the house takes aim to shoot him, preemptively, believing him to be the murderer.

I will reserve the term "Epistemically Justified Mistaken Threat" for Innocent Threats who intend to kill someone whom they believe to be liable to be killed. The resident belongs in this category. He intends to kill someone who is in fact innocent in every sense, so his action is objectively wrong. But given his beliefs, which we may assume are epistemically fully justified, his action is subjectively not only permissible but justified.

(2) *The Technician* A technician is guiding a pilotless drone aircraft toward its landing when it unaccountably veers off course in the direction of a group of houses in which he reasonably believes several families are living. He alters the drone's course in the only way he can, sending it where he knows it will kill one innocent bystander when it crashes. Although there was no reason for him to know this, all the families in fact moved out the day before, while the technician was still on vacation.

Because the plane would not have killed or injured anyone had he not acted, the technician's altering its course is objectively wrong. But his action is subjectively justified because his belief that the plane would otherwise have killed numerous people is fully epistemically justified and, if that belief were true, his act would be objectively justified. So just as the case of the resident is a case in which a person is subjectively justified in intentionally killing a person who is in fact innocent, so this

is a case in which the technician is subjectively justified in foreseeably but unintentionally killing an innocent person.

(3) *The Conscientious Driver* A person who always keeps her car well maintained and always drives carefully and alertly decides to drive to the cinema. On the way, a freak event that she could not have anticipated occurs that causes her car to veer out of control in the direction of a pedestrian.

I will assume that on an objective account of permissibility, this conscientious driver is acting impermissibly. It is impermissible to drive, or to continue to drive, when one will lose control of the car and threaten the life of an innocent person. But of course she cannot know that these are the conditions in which she is driving. So while on an objective account of permissibility, she is an Excused Threat, on a subjective account, she is an Innocent Threat. She does not intend to harm anyone and cannot foresee that she will harm anyone, but she knows that driving is an activity that has a very tiny risk of causing great harm—so tiny that the activity, considered as a *type* of activity, is entirely permissible. But she has bad luck. Notice that although her action is subjectively permissible, it is not subjectively justified. She has no positive moral reason to engage in the activity that she knows has a tiny risk of unintentionally killing an innocent bystander.

(4) *The Ambulance Driver* An Emergency Medical Technician is driving an ambulance to the site of an accident to carry one of the victims to the hospital. She is driving conscientiously and alertly but a freak event occurs that causes the ambulance to veer uncontrollably toward a pedestrian.

This example is just like the case of the conscientious driver except that the ambulance driver has a positive moral reason to drive. She therefore has a subjective justification, and not merely a subjective permission, to act in a way that she knows has a tiny risk of causing great harm to an innocent bystander.

(5) *The Cell Phone Operator* A man's cell phone has, without his knowledge, been reprogrammed so that when he next presses the "send" button, the phone will send a signal that will detonate a bomb that will then kill an innocent person.

It is objectively wrong for the cell phone operator to press the "send" button. But he cannot know that. He is fully epistemically justified in

believing, and with a degree of credence approaching certainty, that his pressing the button is entirely harmless. So if he presses it, he will be acting subjectively permissibly, though not with subjective justification, unless there is a positive moral reason to press it (such as the need to call the emergency number to report an accident). He is therefore an Excused Threat on an objective account, and an Innocent Threat on a subjective account. But unlike the resident, the cell phone operator does not choose to cause harm. And unlike the conscientious driver and the ambulance driver, he does not choose to engage in an activity that has a foreseeable risk of causing serious harm. Subjectively, his situation is not relevantly different from my situation right now, or yours. For after all, my computer could be rigged to detonate a bomb when I press the "tab" key. And even your book could be rigged to detonate a bomb when you turn the next page.

Because anyone who is an Innocent Threat on a subjective account is an Excused Threat on an objective account, while not all who are Excused Threats on an objective account are Innocent Threats on a subjective account, I will refer to the Threats in the five cases cited above as Innocent Threats rather than as Excused Threats. This makes it clear that the reason they are excused on an objective account and acting permissibly on a subjective account is that they are acting on the basis of epistemically justified but false beliefs. Although none of the five is blamable for posing an objectively unjustified threat, they nevertheless differ in the degree to which they are responsible for the threat they pose. Of the five, the one who bears the greatest responsibility for the threat she poses is the conscientious driver. Although she does not intend to harm anyone, she does know that her action carries a small risk of causing great though unintended harm. Although her act is of a type that is generally objectively permissible, and although she has taken due care to avoid harming anyone, she has had bad luck: the risk she knew her act carried has now, improbably and through no fault of her own, been realized. Because she knew of the small risk to others that her driving would impose, and because she nonetheless voluntarily chose to drive when there was no moral reason for her to do so—in short, because she knowingly imposed this risk for the sake of her own interests—she is morally liable to defensive action to prevent her from killing an innocent bystander.

In contrast to the conscientious driver, the resident, the technician, and the ambulance driver all act not merely subjectively permissibly but with subjective moral justification. Each justifiably believes that

he or she has a strong moral reason to do exactly what he or she is doing, and that this reason is not outweighed by any countervailing reasons. But in each case, one or the other of these beliefs is mistaken. Although all of these three agents are blameless, it is reasonable to suppose that all are nevertheless responsible for their choices and that this responsibility, however minimal, is a basis for liability to defensive action. There are, however, differences among them that arguably make a difference to the degree of their responsibility. All three pose a lethal and objectively unjustified threat to a person who is in fact innocent. The resident threatens to kill a person intentionally, the technician threatens to kill a person foreseeably but unintentionally, and the ambulance driver took a risk of killing a person and now as a result of bad luck threatens to kill someone accidentally. If it is true, in general, that it is more seriously wrong to kill a person intentionally than to kill a person foreseeably but unintentionally, and more seriously wrong to kill a person foreseeably than to take a known risk of killing a person accidentally, then the resident chooses to take a greater moral risk in acting than the technician does, and the technician takes a greater moral risk than the ambulance driver. To choose to kill a person intentionally is to take a great moral risk. If one has bad luck and gets it wrong, so that one ends up having intentionally killed a person who is in fact innocent, the degree to which one is responsible for the death seems greater than it is if one has merely run a low risk of killing an innocent person accidentally and ended up killing her through bad luck. If moral responsibility for an objectively unjust threat is the criterion of liability to defensive action, and if the degree of a person's liability varies with the degree of her responsibility, then it seems that the resident is liable to defensive action to a greater degree than the technician, who is in turn liable to a greater degree than the ambulance driver.

4.1.5 Nonresponsible Threats

Like the conscientious driver, the cell phone operator acts in a way that is merely subjectively permissible rather than subjectively justified. There is no positive moral reason for him to press the "send" button on his phone. Yet intuitively he seems even less responsible for the threat he poses than the three Innocent Threats who act with subjective justification. The reason why this is so is that, unlike the others, he does not intentionally kill, knowingly kill, or even knowingly impose on others a risk of being killed. The threat he poses is not one that he