

Part II

Normative Foundations: Law Enforcement or War?

Chapter 5

The Morality and Law of Targeted Killing

Jeff McMahan

In announcing that Osama bin Laden had been killed, Barack Obama declared that “justice has been done.” In saying this, he was implying, or perhaps even asserting, that the justification for the killing was a matter of retributive justice—that is, of punishment. The announcement was immediately followed by celebrations in the streets throughout the United States. These effusions were not expressions of relief at the passing of a grave danger but exultations over the achievement of vengeance against a hated enemy. This understanding of the justification for the killing was largely unchallenged in popular domestic discourse. About a week after the killing, when I suggested during an interview on Wisconsin Public Radio that the killing ought to have been regarded as an act of defense rather than punishment, the announcer remarked that “you may be the first person I’ve heard describe this as a defensive action, [to say] that we did this for defense.”¹

There have actually been attempts to defend the moral permissibility of targeted killing on the ground that it offers both vengeance, satisfying the desire of victims for revenge, and retribution, or the infliction of harm according to desert.² It is obvious, though, that a policy of targeted killing—by which I mean not political assassination generally but the killing of suspected terrorists by agents of the state—cannot be justified by appeal to vengeance or retribution. Some philosophers argue that no one can deserve to be harmed.³ Others argue for the more limited claim that no one can deserve to die, or to be killed. But even if some wrongdoers deserve to be killed, the importance of giving them what they deserve is, on its own, insufficient to justify the risks that a policy of targeted killing imposes on innocent people—most notably, the risk of misidentifying the intended victim and the risk of harming or killing innocent bystanders as a side effect. This becomes particularly clear when one takes into account that retribution alone cannot justify the preventive killing of a person who will otherwise perpetrate an act of terrorism in the near future but has not yet harmed any innocent person. Retribution can justify the killing only of those who have already engaged in terrorism and, it might be thought, does so even when killing them would do nothing to protect innocent people. If pure retribution were our goal, our means would

¹ The interview is accessible at <http://wpr.org/wcast/download-mp3-request.cfm?mp3file=dun110509e.mp3&iNoteID=97290>.

² Steven R. David, “Israel’s Policy of Targeted Killing,” *Ethics and International Affairs* 17 (2003) 111–26.

³ Derek Parfit, *On What Matters*, vol. 1 (Oxford: Oxford University Press, 2011) section 39.

therefore have to be to capture suspected terrorists and try them in court. Only then might we be justified in punishing those found guilty in accordance with their desert. Pure retribution is insufficiently important to justify other means that involve a higher risk of killing innocent people. Indeed, as opponents of capital punishment have plausibly argued, the importance of retribution alone is insufficient to justify the risks involved in killing people even with the safeguards against mistake provided by a criminal trial.

That targeted killing can be justified, if at all, only on grounds of defense is compatible with its being a legitimate means of law enforcement. One might, indeed, argue that targeted killing can be justified as a form of punishment, on the assumption that the principal function of punishment is social defense. While many moral and legal theorists continue to conceive of punishment and defense as entirely distinct, others have recently sought to derive an account of permissible punishment from the principles that govern the permissibility of self- and other-defense.⁴ And most people recognize that at least one legitimate function of punishment is to protect innocent people from those who have demonstrated through criminal action that they are potentially dangerous. Yet it would be a mistake to claim that targeted killing could itself constitute a morally or legally permissible form of punishment. That does not, however, exclude its having a legitimate role in law enforcement. I will return to these matters later.

In considering whether targeted killing can be justified, one must separate the question whether it can ever be morally permissible from the question whether it should be permitted in domestic and international law. These questions are interrelated in complex ways, but at least certain dimensions of each can be considered in isolation from the other. I will address the moral question first and then consider what the legal status of targeted killing ought to be.

I. Morality

There are two basic forms of moral justification that might apply to targeted killing. One appeals to the claim that the potential victim has made himself morally *liable* to be killed by virtue of his moral responsibility for wrongful harm, or a threat of wrongful harm, to others. This claim entails the further claim that he has forfeited his right not to be killed, at least for certain reasons and by certain persons. In this respect, being liable to be killed is like deserving to be killed. The main difference between liability and desert is that the reason given by liability is conditional on the act of killing's being a means or unavoidable side effect of bringing about some good effect, usually the prevention or correction of a violation

⁴ For the view that defense and punishment are distinct, see George P. Fletcher, "Punishment and Self-Defense," *Law and Philosophy* 8 (1989) 201–15, 201; also his "Self-Defense as a Justification for Punishment," *Cardozo Law Review* 12 (1990–1991) 859–66. For accounts that seek to defend punishment by appeal to the permissibility of self- and other-defense, see Thomas Hurka, "Rights and Capital Punishment," *Dialogue* 21 (1982) 647–60; Warren S. Quinn, "The Right to Threaten and the Right to Punish," *Philosophy and Public Affairs* 14 (1985) 327–73; Daniel M. Farrell, "The Justification of Deterrent Violence," *Ethics*, 100 (1990) 301–17; Phillip Montague, *Punishment as Societal Defense* (Lanham, MD: Rowman & Littlefield, 1995); and Victor Tadros, *The Ends of Harm: The Moral Foundations of Criminal Law* (Oxford: Oxford University Press, 2011).

of rights. By contrast, the justification for the infliction of deserved harm is not conditional in this way.⁵

The other basic form of moral justification that might apply to targeted killing is a necessity justification, according to which it can be morally justifiable to kill a person who is not liable to be killed if that is necessary to avoid harms to other innocent people that would be significantly worse. Such a person retains his right not to be killed but the right is, in the circumstances, overridden. Necessity justifications are divided between those that are impartial, or agent-neutral, and those that are agent-relative. The impartial form of necessity justification is often called a “lesser evil” justification and is the less controversial of the two. It asserts that the killing of an innocent person or persons can be morally justified when that is necessary to avert harms to other innocent people that would be substantially greater, *impartially considered*. Note that although such a necessity justification is concerned with consequences, it is not a consequentialist justification. It presupposes that there is a constraint against the killing of an innocent person and denies that it can be overridden whenever the overall consequences of doing so would be better. It requires instead that they be *substantially* better. It is usually held, moreover, that in order to justify the *intentional* killing of an innocent person, the harms that one would thereby avert must be even greater than those whose prevention would be necessary to justify the foreseen but unintended killing of the same innocent person.

The agent-relative form of necessity justification does not require the impartial evaluation of consequences but permits agents to take into account their relations to others. Some philosophers argue that, if one person is related to another in an especially morally significant way—for example, if the one is the parent of the other—there can be a necessity justification for the parent to protect the child by inflicting a harm on another innocent person that is only slightly less than the harm that the child is thereby prevented from suffering. Indeed, some philosophers argue that the parent can have a necessity justification for inflicting a harm on an innocent bystander that is *greater* than that which the child would otherwise have suffered. Such views would have to be considered in any comprehensive discussion of targeted killing, as they suggest the possibility that an instance of targeted killing could be justified even if the harm it would prevent the innocent members of one group from suffering would be *less* than the harm it would cause to innocent members of another group as a side effect. But I will not explore these complications here.

I will, indeed, say little even about the lesser evil form of necessity justification. This is because in those instances in which it is most plausible to suppose that targeted killing is morally justified, such as the killing of bin Laden, the justification seems to be a liability justification. By his own wrongful action, bin Laden had forfeited his right not to be killed if killing him was the best means of preventing innocent people from becoming victims of his terrorist activities. A liability justification is not, however, always decisive. It is possible that there were reasons not to kill bin Laden that made it morally wrong to kill him. But that the killing *wronged* him, or violated his rights, is not among them.

It is perhaps worth mentioning, if only parenthetically, that liability and necessity justifications are in principle combinable. Suppose, for example, that a person, P, has made himself liable to suffer harm up to amount *x* as a means of preventing an innocent person

⁵ This claim has recently been forcefully challenged in John Gardner and François Tanguay-Renaud, “Desert and Avoidability in Self-Defense,” *Ethics* 122 (2011).

from suffering a harm for which P would be partly responsible. Yet to prevent this other harm it is necessary to inflict on P a harm greater than x —say, $x + y$. If the harm that the innocent person would otherwise suffer is sufficiently serious, it could be justifiable to inflict a harm in the amount of $x + y$ on P. The harm that P would suffer up to x would be justified as a matter of liability, while the additional harm, y , would be justified on grounds of necessity. Even though P would be liable to be harmed, the infliction of harm beyond that to which he was liable would have to be justified by reference to the demanding standards that govern the intentional harming of innocent people.

There is, however, a feature of targeted killing that would appear to make it difficult to justify on grounds on liability. This is that targeted killing is *preventive*—that is, it is done not when the victim is engaged in terrorist activity but at a time when he is not attacking, or actively posing a threat. This is a defining rather than contingent feature of targeted killing. The killing of a terrorist while he is attempting to carry out a terrorist attack is not an instance of targeted killing but a straightforward instance of third party defense of innocent people and as such raises no special issues. Only an absolute pacifist might object to the killing of a terrorist as a necessary and proportionate means of thwarting a terrorist attack that is in progress. But how can a person be liable to be killed as a matter of defense at a time when he is not actively posing a threat?

The answer is that a person can make himself liable to be killed if he acts in a way that increases the objective probability that he will wrongly kill an innocent person. For example, a person who plans and prepares for the murder of an innocent person thereby increases the potential victim's risk of being murdered. If the *only* opportunity to prevent the murder occurs in advance of the time that the potential murderer plans to commit the murder, he can be liable to be killed at that time. For even at that time he has made it the case through his own wrongful action that either he must be killed or his intended victim must remain at high risk of being murdered by him.⁶ It is, of course, not certain at the time that, if he is not preventively killed, the potential murderer will later kill his intended victim. Perhaps he will change his mind. But unless the objective probability that he will kill his intended victim is so low that killing him defensively would be disproportionate, it would be unjust for his wholly innocent potential victim to have to bear a risk of being murdered by him in order that he should be spared.⁷

The targeted killing of a person who is in fact a terrorist is morally—though not legally—quite similar to the killing of an “unjust combatant” (that is, a combatant fighting in a war that lacks a just cause) while he is asleep, which most people regard as permissible. A sleeping unjust combatant in a time of war has committed his will to the killing of opposing “just combatants” (who fight in a just war). He intends, or intends conditionally on receipt of an order, to kill them. The broad contours of his life are shaped and guided by this

⁶ For a hypothetical example and further discussion, see Jeff McMahan, “Preventive War and the Killing of the Innocent” in David Rodin and Richard Sorabji (eds.), *The Ethics of War: Shared Problems in Different Traditions* (Aldershot, UK: Ashgate Publishing, 2005) 169–90.

⁷ For further discussion, in which I suggest that even mental acts such as the formation of an intention or even mere deliberation could in principle be a basis of liability to defensive harm, see Jeff McMahan, “The Conditions of Liability to Preventive Attack” in Deen K. Chatterjee (ed.), *Gathering Threats: The Ethics of Preventive War* (Cambridge: Cambridge University Press, 2011).

commitment: he has trained and planned and prepared for this. He is where he is, doing what he does day after day, in order to contribute to his state's unjust war. Much the same is true of the terrorist: he is committed to and guided by the aim of killing innocent people. Both he and the unjust combatant have acted in ways that have raised the objective probability that people who are not liable to be killed (which in my view includes just combatants who fight by permissible means) will be wrongly killed. The main difference between a terrorist who is preparing for his mission or awaiting orders and a sleeping unjust combatant is that the latter keeps about him the visible indicators of his commitment to attack his adversaries, such as his uniform and weapons, while the terrorist seeks to conceal his intentions, preparations, weapons, and identity as a terrorist.

It does not matter to the sleeping unjust combatant's liability to defensive killing whether he has killed in the past. The newly-arrived soldier who has not yet participated in combat is no less liable than the veteran of many campaigns sleeping next to him. The liability of each to defensive action is based on the threat he will pose when he wakes—or, in an extended sense, the threat he poses now—not on what he has done in the past. The same is true of the terrorist. Two people who are together planning and preparing to carry out a terrorist attack may be equally liable to be preventively killed, even if one has conducted such attacks in the past while the other has not. Their liability to defensive action is based on their responsibility for the threat that the defensive action would be intended to prevent, not on their responsibility for unrelated threats from the past.

Whether a person has engaged in terrorism in the past is not, however, irrelevant to the justification for a particular instance of targeted killing. Its primary significance is evidential. If a person is known to have engaged in terrorist activity in the past, that provides some reinforcement for whatever other evidence there is that he is preparing to do so again. There is in general, therefore, less moral risk involved in the targeted killing of a person who has a confirmed history of terrorist activity.

That a person has engaged in terrorism in the past is also relevant to the weight that it is reasonable to attribute to the possibility of mistake. Compare two targeted killings, each of which is based on a mistake. In the first case, there was no reason to believe that the person killed had engaged in terrorism in the past and in fact he had not. He was believed, however, to be preparing to engage in terrorism. But that belief was false: he was not and would never have been involved with terrorism in any way. In the second case, the person killed was correctly believed to have conducted terrorist attacks in the past. It was also believed that he was preparing for another attack, but in fact his career as a terrorist had ended. By the time he was killed he had become entirely harmless.

Neither of these people was liable to defensive killing, as neither posed a threat. Yet the wrong done to the second person, who had been guilty of terrorist action in the past, is less. Because of his history of terrorist action, he is morally responsible for *appearing* to pose a threat of wrongful harm to innocent people. If he had not killed people in the past, or had surrendered himself earlier, other people would not now be forced to choose between killing him and allowing him to live when they reasonably believe that the latter alternative would allow innocent people to remain at risk of being murdered by him. Through his past action, he has forfeited any claim to the benefit of the doubt. He has also, it seems, forfeited his right to kill in self-defense, despite the fact that he no longer poses a threat. He is, one might say, liable to be killed on the basis of a mistake that he is responsible for making it reasonable for others to make, even though he is not liable to be killed for defensive reasons. This is similar

to the claim, which I also accept, that a person can be liable to be killed as a side effect of defensive action even when he is not liable to be killed intentionally as a means of defense.⁸

Another possible reason why the erstwhile terrorist has been wronged to a lesser degree than the person who was innocent of any involvement with terrorism is that he may have deserved to suffer some degree of harm because of his past action. As I noted earlier, it may be doubtful that he deserved to be killed. But if he deserved to be harmed to *some* extent for the harms he inflicted on innocent people in the past, it seems that the *undeserved* harm he suffered in being killed must be less than that of the wholly innocent person. (There is disagreement about whether a person who deserves to be harmed gets what he deserves when he is harmed by natural causes or for reasons unrelated to his desert. Those who think he does not will join those who do not believe in desert in rejecting this second possible reason for thinking that the former terrorist suffers a lesser wrong in being killed than the innocent person does.)

In summary, although targeted killing is necessarily preventive, that does not exclude the possibility of there being a liability justification for it, since people can make themselves liable to be preventively killed. The conditions in which there might be a liability justification for the targeted killing of a terrorist are that, by intending, planning, or preparing to commit or contribute to an act of terrorism, this person is morally responsible for an increase in the objective probability that innocent people will be murdered; that killing him is the best means of averting the threat he poses (both because of the probability of success and because of the expected effects that other options would have on innocent people, including innocent bystanders and anti-terrorist agents); and that killing him is proportionate in the sense that the expected saving of the lives of innocent people substantially outweighs any expected harms that the killing might cause to innocent bystanders as a side effect.

II. Law

That targeted killing can in some cases be morally justified on grounds of liability does not entail that it ought to be legally permitted in those cases. Legal permissions and prohibitions cannot simply restate moral permissions and prohibitions. Although perfect congruence between criminal law and morality is perhaps the ideal, laws must be evaluated on the basis of their likely effects. This may be particularly true of laws governing the action of states, since the abuse of legal permissions by states can have unusually bad consequences. One question, therefore, is whether at least some instances of targeted killing ought to be legal under international law.

But a different and more urgent question is how targeted killing ought to be regarded in relation to the law as it is now. Targeted killing might be thought to come within the scope of either of two legal paradigms. One of these is the set of legal norms governing law enforcement, or police action. The other is the set of legal norms governing the conduct of war. If terrorists are criminals, or criminal suspects, their treatment ought to be governed by the norms of law enforcement. If they are combatants, their treatment ought to be governed by the laws of war. It cannot be the case that terrorists—that is, actual terrorists and not

⁸ For further discussion of the relevance of responsibility for appearances, see Jeff McMahan, “Who is Morally Liable to be Killed in War,” *Analysis* 71 (2011) 544–59. On liability to unintended harm, see Jeff McMahan, *Killing in War* (Oxford: Clarendon Press, 2009) 218–21.

merely suspected terrorists—are neither criminals nor combatants; for if they are not combatants, they are definitely criminals. Yet there may be no determinate, objective truth about which they are as a matter of law. There is certainly no agreement, no consensus, on this matter. There is, it seems, some legitimate scope for choice. The relevant question may not be whether terrorists *are* criminals or combatants but whether it is *better* to classify them as criminals or as combatants.

Whether terrorists are best treated as criminals or combatants, and thus whether anti-terrorist activity is best understood as law enforcement or war, is highly relevant to the status of targeted killing in the law. In the law enforcement paradigm, those who are in fact criminals must be treated as criminal suspects prior to conviction. Their treatment is governed by a requirement of arrest: they must be arrested and tried in a court of law. They may not be hunted and killed, for that would constitute “extrajudicial execution”—a charge often made against targeted killing. If, therefore, terrorists are best regarded as criminals, targeted killing is in most cases illegal. There are exceptions, as I will indicate later, but targeted killing must be ruled out as a *policy* that substitutes for efforts to capture terrorists and place them on trial.

If, by contrast, terrorists are combatants, they may, like other combatants, be permissibly killed at any time during a state of war. The state of war is of course essential for the activation of the laws of war. There is no legal permission for soldiers in one state to kill soldiers in another if the two states are not at war. This is one reason why it was important to members of the Bush Administration to have a “war on terror.” They wanted to kill terrorists as well as to capture them for interrogation; hence they sought to bring their anti-terrorist activities within the scope of the norms governing the practice of war by declaring terrorists to be enemy combatants at war with the United States.

Terrorists often conceive of themselves as combatants and wish to be regarded as such. This may have been part of the reason for Osama bin Laden’s fatuous attempt at a declaration of war against the United States, an act that was merely an attempt because a private person does not have the legal power to declare war. Combatant status has at least two sources of appeal. One is that it confers a specious aura of legitimacy that terrorists sometimes covet. The other is that it might seem to entitle terrorists to the legal rights of prisoners of war when they are captured—rights that, however unrealistically, they sometimes demand. Yet it may actually be against their interests to be recognized as combatants. For that recognition provides their adversaries with a public justification for killing rather than capturing them. The Bush Administration’s violations of rights of *habeas corpus* and its repellent torturing of detainees, along with the Obama Administration’s pusillanimous unwillingness to conduct civilian trials of terrorist suspects, have made the practice of capturing and imprisoning suspected terrorists politically unpopular in United States. The Obama Administration greatly prefers to kill such people rather than capture them—as in the case of bin Laden himself.

One might argue that if terrorists are combatants, that gives them, among their other rights, a right of surrender, which they can use to compel their adversaries to capture rather than kill them. But this ignores two obvious points. First, the targeted killing of suspected terrorists is increasingly done with remotely controlled weapons, such as Predator drones. This denies the victims any option of surrender, which many members of the Obama Administration no doubt regard as an advantage. Second, even when there are opportunities to capture terrorists rather than kill them, there may be little incentive for anti-terrorist agents or their leaders to avail themselves of those opportunities. It may well be that the right of surrender is, unlike some of the rights of prisoners of war, more than merely conventional.

But the motivation to respect it often comes from an expectation of reciprocity. That is, persons on one side of a conflict will be willing to accept the burdens involved in holding prisoners only if they can expect that they and others on their side will, if the opportunity arises, be taken prisoner rather than killed. But terrorists are not in the business of taking prisoners. They “fight” against civilians, not anti-terrorist agents. They take only hostages, not prisoners. There is therefore no basis for an expectation of reciprocity, and thus little reason to expect that terrorist suspects will be offered an opportunity to surrender as long as killing them is politically more expedient.

Although the Bush Administration claimed that terrorists are combatants, it was unwilling to accord them any of the rights that go with combatant status. It therefore declared them to be “unlawful combatants,” a category whose members supposedly have all the liabilities of combatant status, such as being liable to be killed at any time, but none of the corresponding rights or immunities, nor even any of the rights of criminal suspects. The notion of an unlawful combatant is, however, of disputed application. It was originally invoked in the *Quirin* case during the Second World War to justify the execution of a group of German military personnel who had entered the United States clandestinely and were impersonating civilians in an effort to sabotage war-making facilities on American soil. They were official agents of an enemy state who, disguised as civilians, were carrying out military functions in a legally recognized war against the United States. They were not terrorists attempting to kill civilians.

Just as it is unclear what the criteria are for being an unlawful combatant, so it is unclear what rights and liabilities unlawful combatants would have if they could be reliably identified. Certainly their legal status is not what the Bush Administration in practice took it to be—that is, people who may be either hunted and killed or captured and imprisoned indefinitely with no right to legal representation, no right to trial, no right against torture, indeed no rights at all.

The idea that terrorists who are not members of any regular, legally recognized military organization can have some form of combatant status is doubtfully coherent.⁹ Combatant status is a legal artifact. The role of the combatant is defined by reference to legal rights and duties and has been designed so that conferral of combatants status will serve certain purposes—primarily the reduction of violence and harm in war through the insulation of ordinary civilian life from the destructive and disruptive effects of war. The granting of combatant status involves a tacit bargain. Those to whom it is granted are thereby guaranteed immunity from legal prosecution for acts, such as killing and maiming, that would ordinarily be criminal, even if the war in which they fight is unjust and illegal. And they are also granted legal rights to humane treatment and release at the end of the war if they are captured. In exchange for these rights and immunities, they acquire certain duties: they must visually identify themselves as combatants and carry their weapons openly. More importantly, combatants have a legal duty not to conduct intentional attacks against civilians. Combatant status is conditional on reciprocity: one is entitled to the benefits only if one fulfills the duties. Combatants who intentionally kill civilians forfeit some of the privileges and immunities conferred by combatant status—though they do not forfeit combatant status altogether, since even war criminals retain the legal right to kill enemy combatants until they

⁹ For related discussion, see Jeff McMahan, “War, Terrorism, and the ‘War on Terror’” in Christopher Miller (ed.), *War on Terror: The Oxford Amnesty Lectures 2006* (Manchester: Manchester University Press, 2009) 166-70.

cease to be combatants, either when war ends or they are rendered *hors de combat*. (It is one of the many implausible elements of the law of war and the traditional theory of the just war that they permit all combatants, including those fighting for unjust ends, to kill enemy combatants even when the latter are trying to stop them from committing an atrocity.)

While combatant status is thus awarded in part to draw a sharp moral and legal line between those who have it and those who do not, terrorism seeks to erase that line. It is a defining characteristic of terrorism that its instrumental purpose is precisely to expose ordinary civilian life to the violence characteristic of war. Terrorists also subvert the purpose of distinguishing between combatants and noncombatants by concealing themselves among ordinary people and carrying out their attacks without identifying themselves as threateners, thereby limiting the ability of their opponents to distinguish between those who threaten them and those who do not. It is thus the essence of terrorism that terrorists do exactly what the legal category of the combatant has been designed to prevent people from doing. Combatant status is, in effect, a reward offered as an incentive not to do precisely what terrorists do. It would therefore be pointless to grant the rewards for refraining from engaging in terrorism to terrorists themselves.

Despite this argument, many people will remain convinced that terrorists must count as combatants because the dangers they pose often require a military response, as in the case of bin Laden, who had to be killed by a team of elite military commandos. But these people would do well to consider what this idea implies. It implies, for example, that if, on September 11, 2001, members of Al-Qaeda had had a jet of their own that was not intended to resemble a civilian jetliner, and if there had been no one on board other than themselves, their flying it into the Pentagon would have been a legitimate act of war. For the Pentagon is a military headquarters and is thus a legitimate target for enemy combatants during a state of war. One might object that there was no state of war between Al-Qaeda and the United States at that time, but the attack itself would have initiated such a state if the Al-Qaeda operatives had been combatants.

I should clarify that I do not deny that some terrorists can be combatants. But this is not because terrorists generally are combatants but because a combatant can become a terrorist by using terrorist means rather than legitimate military means in an effort to achieve his ends. During the Second World War, for example, political leaders, military commanders, and flight crews collaborated in the bombing of cities with the intention of killing their civilian inhabitants as a means of breaking the morale of their enemies and coercing the enemy government to surrender. These people were engaged in terrorism, which can be deployed in service of just as well as unjust ends. We do not, however, usually refer to such people as terrorists, partly for patriotic reasons if they were on our side, but also because we have another label for regular combatants who commit acts of terrorism: war criminals. Robert McNamara, who was involved in planning the bombings of Japanese cities, made the following observation during an interview conducted late in his life: “Was there a rule that said you shouldn’t bomb, kill, shouldn’t burn to death 100,000 civilians in a night? [General Curtis] LeMay said that if we’d lost the war, we’d all have been prosecuted as war criminals. And I think he’s right. He, and I’d say I, were behaving as war criminals.”¹⁰ He could with

¹⁰ The interview is in a film called “The Fog of War,” directed by Errol Morris, which can be accessed at <http://video.google.com/videoplay?docid=-8653788864462752804#>. The relevant comment occurs about 42 minutes into the documentary. I am grateful to Robert Van Gulick for the reference.

equal justice have confessed that they were acting as terrorists. I will not, however, discuss combatants who become war criminals by engaging in terrorism; rather, in what follows, I will use “terrorist” to refer only to those who engage in terrorism outside of any legally recognized role within a regular military organization. (Some writers tendentiously define “terrorism” so that it can be perpetrated only by “non-state actors.” That is not my suggestion. I am simply limiting the scope of this discussion.)

As I noted earlier, if terrorists are not combatants, they must be criminals. They are civilians who are engaged in an egregious form of criminal activity. Anti-terrorist action is therefore a form of law enforcement, and thus comes within the scope of the norms governing police action. If this is right, terrorists may not be hunted and killed but must instead be arrested and brought to trial.

Critics of this view sometimes object that if anti-terrorism is a form of law enforcement, its aim must be punishment, for the aim of law enforcement is criminal justice—that is, the punishment of the guilty. But anti-terrorism does not aim at punishment; it is, as I claimed earlier, a form of *defense*.

These critics are right that it can be important to keep defense and punishment distinct, even though they are closely related. Although some of the classical just war theorists held that the sole just cause for war is the punishment of the guilty, almost no one holds that view now. Until quite recently, many just war theorists have held instead that the only just cause for war is national defense, either self-defense or third party defense of another state. Retribution, they have held, has no role in the justification of war. According to this view, the reason it is permissible to kill combatants is not that they are guilty and deserve punishment but because killing them is necessary to defend other people from the threat they pose. But if this is right, the idea that it was permissible to kill bin Laden because he was a combatant in the “war on terror” is doubtfully compatible with the idea that his having been killed meant that justice had been done, as Obama proclaimed. To claim both that he could be killed because he was a combatant and that killing him was just punishment is to conflate defense and retribution.

That said, it is important to note both that punishment is only one aim of law enforcement and that one of the functions of punishment is societal defense—that is, the removal of dangerous criminals from society for a period in part to protect the other members of the society from them. But law enforcement also has the protection of innocent people as an aim that is independent of punishment. It can be a legitimate police function to kill a violently dangerous person if he cannot be otherwise subdued for arrest, even if this person is known not to be responsible for his action and thus not someone who deserves to be punished as a matter of retribution.

Defense is an aim of law enforcement in both these ways. When the law aims at defense through punishment, the immediate danger from the criminal has usually passed. A crime has been committed. There is often a threat of further criminal action by the same person, but the need for defensive action may not be urgent. And in most cases of domestic criminal activity, it is normally just as effective and no riskier to law enforcement agents to seek to arrest the suspect than to kill him. Once he has been arrested and no longer poses an immediate threat, it is necessary to try him in court in order to ensure, to the greatest degree possible, not only that no harm is inflicted on an innocent person but also that any harm that is inflicted will be *effectively* defensive—which it will not be if an innocent person is punished by mistake, for in that case the real culprit is left free to cause further harm. The

requirement of arrest is thus both a safeguard against mistake and an important element in the process of ensuring that defensive action is effective.

The second way in which law enforcement can be defensive is quite different. When a criminal suspect evades arrest and poses a clear danger to innocent people, the urgency of defensive action is considerable. Continued efforts to arrest him may leave innocent people—further intended victims, innocent bystanders, and police officers—exposed to a level of risk so high that the requirement of arrest must be suspended. The conditions in which the requirement is suspended resemble those in which private individuals are permitted to kill in self- or other-defense. In such conditions, police officers may then permissibly kill the suspect. Granting law enforcement agents this permission involves significant risks: they may kill the wrong person, they may fail to see that there is an effective alternative to killing, their action may pose a threat to innocent bystanders that is at least as great as that posed by the suspect, and so on. But sometimes these risks are outweighed by the risks involved in failing to eliminate the threat posed by the suspect.

Because anti-terrorist action is generally preventive in character, there is normally less urgency than there is when a violent individual is on a rampage, and the risk of misidentification is significantly greater. In these conditions, it may be reasonable to subject defensive action to safeguards that are not possible, or would be unduly risky, in the case of more immediate threats posed by readily identifiable threateners. In general, therefore, anti-terrorist action should be constrained in the ways characteristic of law enforcement and the administration of criminal justice—that is, there should be a requirement of arrest, a presumption of innocence, an insistence on proof of guilt beyond reasonable doubt, and so on. Observance of these restrictions may even yield an important benefit as side effect: namely, the divulging of information by the terrorist that facilitates the prevention of terrorist acts by others.

Yet there are various other features that characterize much anti-terrorist action that may make it morally necessary on certain occasions to suspend these requirements. Among these features are that terrorists often live, conspire, and train in a state other than the one that is the target of their terrorist action, and that they are often protected by the government of that state and sheltered by local supporters. When anti-terrorist agents can thus expect to be denied permission to make an arrest and to face resistance if they try, the probability of a successful arrest may be low while the risks involved in the attempt may be high. When an unusually dangerous terrorist is inaccessible to arrest at a reasonable level of risk for these or other reasons, conditions may be analogous to those that justify the suspension of the requirement of arrest in cases of domestic law enforcement. In these conditions, targeted killing may be justified for reasons similar to those that can justify the police in killing a rampaging gunman who resists arrest.

Many of these conditions obtained in the case of Osama bin Laden. He had proven himself to be a highly dangerous terrorist. There was no risk that, in killing him, anti-terrorist agents would be killing an innocent or unthreatening person. And there can be little doubt that he was being sheltered by certain individuals in the Pakistani government or military, or both. Any effort to secure the cooperation of the Pakistani government or military in arresting him would therefore almost certainly have resulted in his being alerted and allowed to escape. Finally, there was good reason to believe that he was heavily protected by armed guards. Yet in spite of all this, it turned out to have been possible to capture him alive at little or no more risk than was involved in killing him. The initial reports revealed that, *before* they shot bin Laden himself, the Seals incapacitated a woman who charged them as they entered bin

Laden's room by shooting her in the leg. That immediately raised the question why they could not have done the same with him. The Obama Administration soon conceded that he was unarmed and, despite the Administration's assertion that the Seals were "prepared" to capture him if possible, *The New Yorker* has quoted "a special-operations office who is deeply familiar with the bin Laden raid" as saying that "there was never any question of detaining or capturing him—it wasn't a split-second decision. No one wanted detainees."¹¹

It seems, then, that according to the view for which I have been arguing, it was wrong to kill bin Laden rather than capture him. The reason has nothing to do with the fact that he was defenseless. That is a distraction, a sentimental relic of medieval codes of chivalry. If killing bin Laden had been necessary to eliminate a significant threat for which he was responsible, even if he did not *pose* that threat at the time, it would have been unambiguously *good* that he was defenseless when the killing had to be done, so that no harm might be done to those acting justifiably to eliminate the threat. The reason is instead that killing him does not seem to have been necessary to avert any threat for which he was responsible.

Yet that might not be true. There are at least two reasons that may have motivated members of the Obama Administration to order that he be killed rather than captured, either of which might provide a justification for the killing, despite the many reasons why it would have been desirable to capture him and place him on trial. One is that the Administration may have decided in advance that it was not worth the loss of even one more American life to enable bin Laden to live to face trial rather than be killed. So the Seals may have been instructed simply to take no chances. The other is that the Administration may have reasonably feared that if he had been taken captive, his followers would then have taken American hostages and begun killing them one by one in an effort, however futile, to coerce the United States to release him. It is not unreasonable to suppose that the reasons favoring capture rather than killing may have been outweighed by that risk. (I would not include among those reasons that he had a right not to be killed. If killing him was necessary to avoid a significant risk that innocent people would be killed by his followers, then he was liable to be killed, as he would have borne some responsibility for the acts of his followers. This is the kind of case to which I referred in the previous paragraph.) The risk that hostages might be taken in an effort to secure his release could, however, have been minimized if the Obama Administration had postponed the announcement of his capture long enough to have placed him in the custody of an international body, such the International Criminal Court in the Hague. But this option was of course politically impossible in the United States, where the outrage and jeers of Republican politicians at the Administration's placing the United States' greatest enemy under international jurisdiction would have converted the capture from a triumph to a humiliation. As recent experience demonstrates, Republican politicians can be counted on to obstruct the best solution to any problem.

There are various reasons why capture followed by trial is generally preferable to killing. Apart from the fact that a dead terrorist can provide no information about other terrorists or planned terrorist operations, most of the disadvantages of targeted killing have to do with the risks it involves that can be mitigated through the safeguards provided by the alternative of capture and trial. Perhaps the most obvious risk is that the victim may be misidentified. In one of the earliest instances of targeted killing, agents of Mossad, the Israeli intelligence and counter-terrorism agency, killed an innocent Moroccan waiter in Norway in 1973 in the mistaken belief that he was the leader of the Palestinian "Black September" group

¹¹ Nicholas Schmidle, "Getting Bin Laden," *The New Yorker* (August 8, 2011) 43.

that had massacred Israeli athletes at the 1972 Munich Olympics. This case provoked an international scandal, but in general the incentives to exercise reasonable care in identifying and attacking foreign terrorists are weaker than those for exercising care in domestic police work. Governments naturally take greater precautions to avoid killing their own citizens by mistake. Another instance of misidentification occurred in London when British police killed a Brazilian man whom they mistook for a terrorist shortly after the terrorist bombings there in 2005.

In addition to the mistake of misidentification, there is also the possibility that killing someone known to have engaged in terrorist action in the past will serve no defensive purpose, perhaps because the person has altogether ceased to be involved in terrorist activity. As I noted earlier, however, the wrong done to the victim in this kind of case is significantly less than it would be if he had not been a terrorist and thus bore no responsibility for the reasonable belief of others that he continued to pose a threat of wrongful harm. For much the same reason, a lesser wrong is also done when a member of a terrorist organization is killed when his contribution to the organization's action was, while sufficient for liability to a lesser form of harm, insufficiently significant to make him liable to be killed.

Another risk of targeted killing that might be lessened by pursuing the alternative of capture and trial is the harming of innocent people as a side effect. Because terrorists tend to live and move freely among ordinary people, it is difficult to attack them without killing or injuring innocent bystanders as well. This is particularly true when targeted killing is attempted using remotely controlled weapons. This problem is not, however, unique to targeted killing; it arises as well for defensive action taken in response to an actual terrorist attack and even for efforts to capture or arrest a terrorist suspect who can be expected to engage in violent resistance. In some cases, indeed, targeted killing can be carried out with almost no danger to innocent bystanders. The classic example is the killing of Hamas's bomb maker, Yahya Ayyash, by agents of the Israeli security service, Shin Bet, who managed to transfer to him a cell phone rigged with explosives, which they then detonated when they confirmed that he was using it. More recently, the targeted killing of Osama bin Laden was accomplished without harm to any innocent bystanders.

In addition to the risks of mistake, there are also risks of abuse. Even in the case of a government that is scrupulous in limiting its use of targeted killing to cases of confirmed terrorists, subtle forms of abuse are likely to develop, such as carelessness about side effects if those involved believe that they can get away with it without adverse publicity. But the most serious form of abuse by a government that kills only confirmed terrorists is one that, as I mentioned earlier, characterizes the Obama Administration's policy of targeted killing. This is the use of targeted killing as a tactic of first rather than last resort, as a replacement for other forms of anti-terrorist action, such as capture and trial, that incorporate stronger safeguards against the inadvertent killing of innocent people.

The greatest danger from any legal recognition of the permissibility of targeted killing is, however, that unscrupulous regimes will exploit that legal permission in offering public justifications for the killing of political opponents who are not terrorists at all but will be said to be by their killers. There is ample precedent for this—for example, the killing in 1982 of the anti-Apartheid activist Ruth First with a parcel bomb sent by agents of the South African government. That government had declared the African National Congress to be a terrorist organization; hence anyone associated with it could conveniently be branded a terrorist. A similar targeted killing was even carried out in Washington, DC, in 1976 when agents of the Pinochet dictatorship in Chile detonated a bomb under the car of Orlando Letelier, a former

minister of the government that Pinochet had overthrown and an opponent of the new regime. In this case the killers were working for a regime that had seized power with U.S. assistance and maintained close ties to the Ford Administration, so protests were muted. But the United States may have to reconsider the precedent it is setting with its targeted killings when regimes with leaders similar to Saddam Hussein, Muammar Qaddafi, or Kim Jung Il acquire small, remotely controlled drones that can be used to kill their opponents on U.S. soil.¹²

The careful development and elaboration of this objection to targeted killing is the main aim of Jeremy Waldron's contribution to this volume and I will not attempt to improve on his superb exposition. It is worth remarking, however, that the objection is compatible with the recognition that targeted killing may in some instances be morally permissible, or even morally required. One might, therefore, accept the same view about targeted killing that some, myself included, have argued is the right view of torture: namely, that while it can on some occasions be morally permissible, it ought to be categorically prohibited by law. According to this view,

if we grant any legal permission to use torture, particularly one that attempts to capture the complex conditions of moral justification, it will be exploited by those whose aims are unjust and [will be] either abused or interpreted overly generously even by those whose aims are just. Throughout human history, torture has been very extensively employed, but the proportion of cases in which the use appears to have been morally justified seems almost negligible. . . . Any legal permission to use torture, however restricted, would make it easier for governments to use torture, and would therefore have terrible effects overall, including more extensive violations of fundamental human rights. The legal prohibition of torture must therefore be absolute. . . . We cannot proceed with torture the way we have with nuclear weapons—that is, by permitting it to ourselves while denying it to others by means of security guarantees, economic rewards, and other measures designed to make abstention in the interests of all. If we permit ourselves to use torture, we thereby forfeit any ability we might otherwise have to prevent its use by others. . . . Our only hope of being able to impose legal and other constraints on the use of torture in the service of unjust ends by vicious and cruel regimes is to deny the option to ourselves as well, even in cases in which we believe it would be permissible.¹³

It may be tempting to argue that, whatever may be true about the law, a state's adversaries cannot make it impermissible for that state to engage in otherwise permissible acts of targeted killing simply because that action would encourage *them* to act impermissibly, or provide a rationale for their doing so. One might say that if a state's

¹² As I was making final revisions to this essay, the Obama Administration accused Iran of plotting to kill the Saudi ambassador to the United States. The Justice Department's accusation can be found at <http://www.justice.gov/opa/pr/2011/October/11-ag-1339.html>. This is not the first time that the theocratic regime in Iran has engaged in targeted killing. See Roya Hakakian, *Assassins of the Turquoise Palace* (New York: Grove Press, 2011).

¹³ Jeff McMahan, "Torture in Principle and in Practice," *Public Affairs Quarterly* 22 (2008) 124–6.

otherwise justified use of targeted killing would prompt unjust regimes to engage in the unjustified use of targeted killing, their unjustified use is not attributable to that state's action and thus cannot make that action disproportionate or otherwise impermissible.

Even if this were true as a matter of morality, however, it would not address the claim that legal arrangements that would permit the targeted killing of, for example, Osama bin Laden by the United States would on balance be worse for everyone, including the United States, because these arrangements would eventually be exploited by all states, not just those that subject the practice of targeted killing to stringent procedural constraints and have the ability to conduct these killings in a reasonably discriminating way. But even as a matter of morality, the view that the permissibility of one agent's action cannot be affected by what it might prompt other agents to do is untenable. To the extent that one can predict what unjust agents would do in response to one's action, one may have to regard their responsive action the way one would regard a natural event that one's action would trigger. If one's action would precipitate an avalanche that would kill a certain number of people, that weighs against the action's being proportionate. Similarly, if one's action would provoke a despot to kill an equal number of innocent people, either intentionally or as a side effect of responsive action, that too may render one's action disproportionate. Even if one's action would otherwise be permissible, the fact that the despot would bear full *responsibility* for the killings is insufficient to render one's action *permissible*.¹⁴

Consider a simplified example. Suppose there are two equally important military targets but we can attack only one of them. If we attack one, the explosion will precipitate an avalanche that will kill 50 innocent bystanders. If we attack the other, our adversaries will kill 51 innocent bystanders they would otherwise not kill. Suppose that either attack would be proportionate. In neither case would we directly kill the innocent bystanders. But in both our action would precipitate and be a necessary condition of the event that would be the proximate cause of their deaths. If we think that killings done by others cannot affect the proportionality of our action, or that effects mediated through the agency of others must be discounted in the determination of proportionality, then we ought to attack the second target, so that 51 innocent bystanders will be killed. People will disagree about this, but my view is that we ought in these circumstances to do what will cause the fewest deaths of innocent bystanders, other things being equal.

A more significant objection to the claim that any legal recognition of the permissibility of targeted killing will be abused by vicious regimes is that the same is true of the alternative means of anti-terrorism consisting of arrest, trial, and punishment. It has always been possible for repressive, nondemocratic regimes to seize their political opponents, subject them to a sham or rigged trial, and then execute them, claiming that justice has been done. This problem is exacerbated for the United States by its embrace of capital punishment. Although the United States has not recently legally executed a foreign terrorist, it did execute a domestic one—Timothy McVeigh—and may eventually execute one or more terrorists captured in the “war on terror” and tried by a military court. The practice of judicial execution in the United States sets a dangerous precedent just as the practice of targeted killing does. Yet even the abolition of capital punishment in the United States would only weaken rather than dispel this concern about the precedent-based objection to targeted killing. For repressive, nondemocratic regimes can also exploit the legal mechanisms of trial

¹⁴ For further discussion, see Jeff McMahan, “Responsibility, Permissibility, and Vicarious Agency” *philosophy and Phenomenological Research* 80 (2010) 673–80.

and imprisonment to silence their political opponents indefinitely, and to deter other potential opponents from engaging in political action.

One could respond to this problem by insisting that norms of anti-terrorist action requiring arrest and trial must specify standards of fairness and openness for trials and sentencing. There would, for example, have to be transparency and public disclosure of the evidence against the suspect. This would prohibit the use by the United States of secret military tribunals or trials that grant fewer rights to defendants than it would demand that its adversaries grant to U.S. citizens in trials they might conduct.

But if one claims that there can be a neutral norm of arrest and trial provided that certain constraints are imposed on what counts as acceptable forms of arrest, trial, and punishment, then a parallel claim might be made on behalf of a neutral norm of targeted killing. Perhaps there could be a neutral norm that permits targeted killing provided that it set high standards of *post facto* justification, with requirements for the disclosure of evidence, a demonstration that killing was both necessary and proportionate, and so on. There could then be legal provisions for international sanctions against states that failed to satisfy the demand for *post facto* justification.

Perhaps, therefore, targeted killing has more in common with ordinary killing in self-defense than we have thought. For there is ample scope for abuse of the legal permission to kill in self-defense in domestic criminal law. If, for example, a woman has a husband with a known record of physical violence, she may be able to provoke him to hit her, then murder him in their home, and afterwards make a successful plea of self-defense at trial. Yet we do not respond to this risk by denying that there can be a neutral rule permitting killing in individual self-defense. We recognize that a trade-off has to be made between the need to permit self-defense in a great range of cases and the need to deter the exploitation of that permission by would-be murderers. But we resolve that problem by imposing a variety of legal constraints on the right of self-defense rather than by denying a legal right of self-defense altogether. Admittedly, these legal constraints can be reasonably effective in domestic criminal law because there are institutions that can enforce them, whereas there are no even remotely comparable enforcement mechanisms in international law. But the difficulty of enforcement is as much a problem for the legal prohibition of targeted killing as it is for the imposition of constraints on a limited legal permission to engage in targeted killing.

The trade-off between the wrongful harms that might be prevented by legally permitting some instances of targeted killing and those that might be facilitated by the exploitation of that permission should be negotiated differently from the trade-off between the wrongful harms that might be prevented by legally permitting some instances of torture and those that might be facilitated by the abuse of that permission. This is mainly because targeted killing may often be both necessary and effective in preventing or limiting terrorist action, whereas torture can rarely be effective as a means of defense. Thus, if there were a limited legal permission to engage in targeted killing, the ratio of justified to unjustified targeted killings would likely be much higher than the ratio of justified to unjustified instances of torture if there were a limited legal permission to engage in torture. It seems, therefore, that the argument cited earlier against even a limited legal permission to practise torture cannot be extrapolated to the case of targeted killing, or not without significant qualification.

Waldron articulates another concern, which is that acceptance of targeted killing will erode the distinction between combatants and noncombatants as it functions in the war convention. The convention of noncombatant immunity has evolved over a long period of time, has a variety of supporting rationales (it limits the violence of war, protects the rights of the innocent, and so on), and is generally believed, though perhaps mistakenly, to have deep foundations in basic, nonconventional morality, so that many people believe that to violate it is to be guilty of murder. It is hard to deny that this convention is of great practical importance.¹⁵ Would the legal acceptance of targeted killing undermine it?

It is at least worth considering whether there could be an effective firewall between the targeted killing of terrorists in peacetime and the killing of civilians in war. Note that not all killings of civilians are legally prohibited in war. It can be legally permissible foreseeably to kill innocent civilians as an unavoidable and proportionate side effect of an attack on a military target. And it can be permissible to kill a civilian if he is armed and threatens the life of a combatant. What is prohibited by the principle of noncombatant immunity is the *intentional, nondefensive* killing of civilians as a means of coercing others, usually their political leaders. What this prohibited form of killing and targeted killing have in common is that they both involve the intentional killing of civilians. But the difference between them is salient and of obvious moral significance: namely, that what is prohibited in war is the intentional killing of people who do not threaten to harm or kill anyone, whereas the victims of morally justified targeted killing are terrorists who will otherwise harm or kill innocent people. Although terrorists may be civilians, they are civilians who have made themselves morally liable to defensive killing. So what the principle of noncombatant immunity prohibits is terrorism, or the use of terrorist tactics in war, while targeted killing aims to prevent acts of terrorism through the killing of terrorists. Rather than subverting noncombatant immunity, therefore, morally justified instances of targeted killing protect innocent, unthreatening people from terrorist attacks. They enforce the principle of noncombatant immunity.

This difference between the killing of civilians for terrorist reasons and the targeted killing of terrorists themselves is sufficiently clear that people everywhere can understand it. Anyone can see the difference between the killing of Ruth First and the killing of Osama bin Laden. There are thus clear criteria by which it could be shown that a legal permission to engage in targeted killing was being abused.

Ideally what is needed is cooperation both among national law enforcement agencies and between national and international law enforcement agencies, so that terrorists can be dealt with efficiently solely by means of traditional methods of law enforcement. But this cannot happen while law enforcement agencies in states that harbor terrorists are controlled by governments that support the terrorists. In these conditions, what is needed is a new body of anti-terrorist law based on the recognition that terrorists are neither combatants in the legal sense nor ordinary criminals but instead have an intermediate status that combines elements of criminality with elements of combatancy. Although they have some of the defining features of both criminals and combatants, terrorists lack some of the defining characteristics of combatants and are considerably more dangerous than ordinary criminals. It is for these reasons that anti-terrorist action cannot be well governed within either the liberal egalitarian paradigm or the war convention. A new body of law designed specifically to

¹⁵ See Jeremy Waldron, "Civilians, Terrorism, and Deadly Serious Conventions" in his *Torture, Terror, and Trade-Offs: Philosophy for the White House* (Oxford: Oxford University Press, 2010).

regulate anti-terrorist action is therefore urgently needed. It is possible that a tightly circumscribed and constrained permission for targeted killing could be a part of that law. But it would have to be formulated to take account of the grave risks to which Waldron calls attention. And given those risks, its acceptance might be explicitly provisional. There is no reason why certain elements of a new law might not be adopted on a trial basis, to be repealed if, once they have been implemented, their costs appear to outweigh their benefits. But what the precise elements of a new body of law designed to govern anti-terrorist action ought to be is obviously a matter that should be addressed to people better qualified than I.