

Professor Phillip Montague

Western Washington University, Department of Philosophy

## DEFENDING DEFENSIVE TARGETED KILLINGS

As is the case with all homicides, there exists a moral presumption against targeted killings. Circumstances can arise, however, in which this presumption is defeated for homicide in general, and – perhaps - for targeted killing in particular. The question to be addressed here is whether the presumption is defeated in the specific circumstances that typically surround certain targeted killings of terrorists that occur as an implementation of official U.S. policy.

A potentially effective way in which to argue for an affirmative answer to this question (and against the idea that targeted killings are nothing but murder) is by demonstrating that the moral permissibility of the targeted killings to which it refers follows from the best available theory of self-defense. There are, however, at least two importantly different ways in which such an argument might be formulated.

On one formulation, a selected theory of self-defense is applied directly to individual targeted killings. The other formulation focuses on national self-defense rather than individual self-defense. Its central claim is that targeted killings are morally permissible as a means by which the United States exercises its right of self-defense against terrorist aggression.

In this essay, I offer reasons for doubting that either of these approaches to defending targeted killings is viable. I also propose an alternative account that I believe

identifies conditions under which targeted killings count as morally permissible defensive homicide.<sup>1</sup> The proposed account applies theories of self-defense to targeted killings in conjunction with an explanation of how individuals can perform actions that are *jointly* defensive or aggressive.

This alternative account allows targeted killings to count as defensive homicides even if they aren't individually defensive, and even if the actions to which they are responses aren't individually aggressive. At the same time, the proposed account eliminates the apparent need for attributions of agency to political communities and terrorist organizations *per se* - to the United States and al-Qaeda in particular. As is explained in the essay, such attributions are at best unhelpful in the present context.

Although my primary purpose is to explain how targeted killings can be defensive and morally permissible, I conclude by considering certain conditions under which they might be impermissible even if defensive.

## I.

Targeted killings can obviously take place in a wide variety of circumstances, although these days they are most closely associated with missile strikes delivered by pilotless drones that are directed at members of al-Qaeda. These are the targeted killings that will be focused on here. For ease of exposition, however, I will ignore certain high-tech aspects of these killings, and assume that each drone is under the control of a single

individual. Accordingly, I will be referring periodically to the following hypothetical case:

AI is assigned the mission of using the drone he controls to assassinate an al-Qaeda leader. This particular terrorist is responsible for fabricating devices that are used by suicide bombers who detonate their explosives in places frequented by large numbers of civilians. AI is provided with reliable information regarding the terrorist's plan to travel by car to his bomb factory. AI launches his drone, locates the terrorist's car, and destroys it with one of the drone's missiles.

Although hypothetical and somewhat artificial, this case possesses the features in virtue of which actual targeted killings are morally controversial. It therefore serves the purposes of this discussion. As is indicated in the introduction, I propose here to identify conditions under which such targeted killings count as morally permissible defensive homicide.

In view of how theories of self-defense are invariably formulated, their primary applications are directly to defensive actions performed by individuals in response to aggressive actions performed by other individuals. Judith Thomson depicts a hypothetical situation of this sort in her classic paper "Self-Defense and Rights":

Suppose Aggressor has got hold of a tank. He has told Victim that if he gets a tank, he's going to get in it and run Victim down. Victim sees Aggressor get in his tank and start towards Victim. It is open country, and Victim can see that there is no place to hide, and nothing he can put between himself and Aggressor which Aggressor cannot circle round. Fortunately, Victim happens to have an anti-tank gun with him, and it is in good working order, so he can use it to

blow up the tank, thereby saving his life, but of course thereby also killing Aggressor.<sup>2</sup>

Assuming that, when Aggressor launches his attack, Victim's actions are morally innocuous, and assuming too that no one else is involved in this situation, all theories of self-defense with which I am familiar would imply that Victim is morally permitted to kill Aggressor.

Of course, our example of Al and the terrorist is very different from Thomson's. The terrorist isn't acting aggressively, and Al's action isn't self-defensive, so the question of whether Al is permitted to kill the terrorist in self-defense doesn't even arise. Theories of self-defense almost invariably apply as well to other-defense, however. For example, a key component of Thomson's account of self-defense is her claim that a person has no right to life if killing him is the only way in which to prevent him from performing an action that would violate another's right to life (and other things are equal).<sup>3</sup> Nothing in this account implies that it is restricted to *self*-defensive homicide. A person who does lack a right to life in virtue of satisfying the condition stipulated by Thomson can be permissibly killed by anyone in a position to do so.<sup>4</sup>

We must therefore consider whether targeted killings like the one carried out by Al could be morally permissible in virtue of being appropriately other-defensive. Taken in context, the following remarks by Andrew Altman and Christopher Heath Wellman suggest that they might answer this last question affirmatively:

Surely, it would have been permissible for someone to have assassinated Stalin in the 1930s. It seems, then, that political assassination is, in principle at least, a morally permissible means of stopping or halting human-rights abuses.<sup>5</sup>

These remarks raise an important question regarding the conditions that govern morally permissible defensive homicide. This question concerns what might be called the “inevitability condition,” according to which killing x in self- or other-defense is morally permissible only if x’s killing an innocent person would otherwise be inevitable.

It might be thought that “highly probable” should replace “inevitable” in the aforementioned condition. But suppose that x is attacking y at time t, and that – at t – x’s killing y is highly probable if y doesn’t kill x first. Suppose too that, as a matter of fact, x’s attack will fail regardless of what y does. Is y nevertheless permitted to kill x at t? I would argue that it isn’t.<sup>6</sup> Probability considerations can, of course, play an epistemic role in the contexts we are examining. However, whether a defensive homicide is permissible depends on whether it is in fact the only way in which to avoid an innocent’s death – not on whether anyone reasonably believes that it is. Such epistemic considerations are (indirectly) relevant to moral culpability or blameworthiness after the fact, but they are irrelevant to moral permissibility before the fact.<sup>7</sup>

The underlying distinction here is between negative moral appraisals of actions and negative moral appraisals of their agents. My primary focus is on the former appraisals – not because they are more important than the latter – but because they are more basic in this respect: a person is morally blameworthy or culpable *for performing a certain action* only if the action is morally impermissible. Blameworthiness and culpability for acting also depend on what the agent could reasonably have been expected to believe about the nature of the action and its consequences; and here is where probability considerations enter into the picture. I will comment briefly on the matter of

blameworthiness for targeted killings after arriving at some conclusions regarding their permissibility.

The preceding remarks lead to a second point that needs to be made before proceeding – namely, that inevitability is distinct from imminence. That inevitability doesn't imply imminence becomes clear on considering the following variant of Thomson's example:

Some Third Party (rather than Victim) has the anti-tank gun. Aggressor isn't yet aware of Victim's presence, but, when he does see Victim, he will launch his attack. Third Party can prevent Aggressor from killing Victim only if he fires his weapon before Aggressor begins his attack.

Third Party's killing Aggressor in these circumstances would be *preemptive* other defense, and would satisfy the inevitability condition referred to above even though the threat posed by Aggressor isn't imminent.<sup>8</sup> Killing Aggressor would also be morally permissible according to some familiar theories of self-defense (including Thomson's own theory).

As it stands, however, our example of Al and the terrorist isn't about preemptive other-defense, since nothing in the example suggests that the terrorist leader will inevitably attack and kill innocent people if he isn't killed.<sup>9</sup> But the example does seem relevantly similar to this variation on Thomson's theme:

Prior to launching his attack on Victim, Aggressor must refuel his tank, and Accessory's truck is the only available source of additional fuel. Accessory is happy to help, since she also wants Victim killed. The only way in which Third Party can prevent Aggressor from eventually attacking Victim and running him

down with his tank is by destroying Accessory's truck and Accessory along with it.

Third Party can save Victim's life by – and only by – killing Accessory. Doing so wouldn't be individually other-defensive, however, since Accessory isn't attacking Victim. It is therefore difficult to see how theories that are aimed at determining the moral dimensions of individually defensive homicide can have anything to say about this case – or our example of Al and the terrorist.

Thomson's theory is a case in point. Its central thesis (one component of which was stated above) can be put as follows: killing *x* is morally permissible if *x* lacks a right to life; and *x* lacks a right to life if killing him is necessary to prevent him from violating another's right to life (and other things are equal). Since it doesn't appear that killing Accessory is necessary to prevent *him* from violating Victim's right to life, Thomson's theory doesn't seem to imply that Third Party is permitted to kill Accessory.<sup>10</sup>

Of course, self- and other-defense situations aren't the only ones in which people are permitted to prevent the deaths of innocents by killing noninnocents. There are also what might be called "self-preservation" and "other-preservation" situations, the latter of which might be thought to include Third Party's killing Accessory and Al's killing the terrorist. By itself, however, this change of direction would be incapable of solving the problem at hand. This is because the permissibility of killing individuals in circumstances like those surrounding Accessory and the terrorist depends on connections between their actions and the actions of others.

Hence, an explanation of the permissibility of Third Party's killing Accessory would need to incorporate an account of how Accessory's and Aggressor's actions

combine to create a lethal threat for Victim. In a parallel fashion, the permissibility of Al's killing the terrorist couldn't be explained without locating the terrorist's actions within a nexus that includes the actions of others associated with Al-Qaeda – suicide bombers in particular. It is worth bearing in mind, by the way, that drones aren't actually operated by single individuals, and that killing the terrorist would require the actions of many individuals serving in various capacities.

These remarks shift the focus of this inquiry from actions that are individually defensive or aggressive to actions that are in some sense collectively defensive or aggressive. This shift in focus leads quite naturally to the idea that targeted killings are matters of national self-defense – that their permissibility follows from the role they play in defensive actions by the United States against aggression by Al-Qaeda. This is the second of the two approaches to the moral problem of targeted killings to which I referred in my introductory remarks.

## II.

On April 14, 1943, a Japanese radio transmission intercepted by U.S. naval intelligence revealed that Admiral Isoroku Yamamoto, the commander-in-chief of the combined Japanese naval fleet, was planning an inspection tour of Japanese military installations in the South Pacific. The transmission also contained a detailed itinerary of Yamamoto's trip, and a decision was made to intercept and kill him. On April 18<sup>th</sup>, American fighter planes located and shot down the plane in which Yamamoto was flying.

If someone were to question the morality of Yamamoto's assassination, the likely response would be that, like ordinary Japanese soldiers, sailors, or airmen, Yamamoto was a combatant waging a war of aggression against the United States, and killing him was therefore morally permissible. Critics of targeted killings commonly receive a similar response: just like members of al-Qaeda who wield weapons and plant bombs, the organization's rear echelon members are also combatants waging a war of aggression; and since killing the former is morally permissible, so is killing the latter.

If al-Qaeda is waging a war of aggression, then the United States is fighting a defensive war. A more precise version of the preceding argument can therefore be formulated as follows:

The United States has a right of self-defense; and establishing and implementing a policy of targeted killings is a means by which the United States defends itself against al-Qaeda aggression.

Hence, the U.S. is exercising its right of self-defense in establishing and implementing a policy of targeted killings. Since implementing this policy requires that targeted killings actually occur, targeted killings are permissible and, in particular, Al is permitted to kill the terrorist leader. Call this "the collective-action argument."

A close examination of the collective-action argument reveals, however, that it is problematic in two respects.

Let us assume for the moment that the United States does indeed have a right of self-defense, and also that engaging in targeted killings is a means by which the United States is defending itself against aggression by al-Qaeda. These two assumptions don't imply that the United States has a right – and is therefore permitted – to engage in

targeted killings. Hence, the conclusion of collective-action argument therefore can't be used to establish the permissibility of targeted killings.

That the collective-action argument is invalid becomes apparent on considering the following parallel argument: Archie has a right of self-defense; killing Betty is a means by which Archie defends himself against Betty's aggression; Archie therefore has a right to kill Betty. This argument is clearly invalid. After all, Archie might be able to defend himself by injuring Betty only slightly, or not at all. Or, putting the point in general terms: the proposition that x has a right to perform acts of a certain type doesn't imply that x has a right to perform every individual action of that type.

The problem with the collective-action argument that we have been examining concerns inferences from very general attributions of particular rights to more specific attributions of those rights. The collective-action argument contains an additional kind of questionable inference, however – namely, the inference from propositions about the United States and al-Qaeda, to propositions about individuals who are related to the United States or to al-Qaeda in certain ways.

This second problem with the collective-action argument arises from the fact that political communities (including the United States) and terrorist organizations (including al-Qaeda) are abstract entities. As such, they can't possibly possess the mental properties such as beliefs and desires (and hence intentions) that are prerequisites for the exercise of agency. Since the United States *per se* and al-Qaeda *per se* are incapable of exercising agency, they can't perform the defensive and aggressive actions to which the collective-action argument refers.<sup>11</sup>

Granted, we commonly use sentences that attribute actions to abstract entities. The use of such sentences is obviously convenient, and the sentences often express true propositions. For example, the sentence “Green Bay played in the Super Bowl” expresses a true proposition if uttered in the proper context, and the same is true of “The Royal Shakespeare Company performed *Romeo and Juliet*.” When these sentences do express true propositions, however, actions and agency are attributed only to concrete individuals who were related in certain ways to the Green Bay Packers in February of 2011, and to other individuals who were members of the Royal Shakespeare Company during certain periods (during March of 2010, for example).

But even though the only actions to which these propositions refer were performed by individuals, football games and dramatic performances aren’t simply collections of individual actions. More specifically, the true proposition expressed by “The Royal Shakespeare Company performed *Romeo and Juliet*” wouldn’t be equivalent to anything like this:  $x_1$  performed  $a_1$  at  $t_1$  and  $x_2$  performed  $a_2$  at  $t_2$  and . . . . In addition to such references to actions performed individually, the true proposition would also contain references to actions that are performed *in concert* or *jointly* with others – references that presuppose the concept of joint agency. Note too that, although the proposition in question concerns a dramatic performance, not all of the actions to which it refers are individual dramatic performances. Some of these references are to actions on the part of individuals who manipulate scenery or control lighting, for example.

In a parallel fashion, if sentences attributing defensive actions to the United States and aggressive actions to al-Qaeda express true propositions, then these propositions refer only to actions that are performed by concrete individuals who are related in certain ways

to the United States and to al-Qaeda respectively. Some of these references are to actions that are performed jointly with others through the exercise of joint agency. Not all of these actions need be individually defensive or aggressive – just as the actions to which the proposition expressed by “The Royal Shakespeare Company performed *Romeo and Juliet*” refers aren’t all individual dramatic performances.

As will be explained presently, the proposition expressed by “The United States is waging a defensive war against aggression by al-Qaeda” might have room for references to targeted killings. Determining whether this is indeed the case requires a theoretical framework within which the concepts of joint action and joint agency are explicable. This theoretical framework provides a means of approaching the moral problem of targeted killings that differs significantly from the two problematic approaches examined in the preceding discussion. This third alternative doesn’t apply theories of self-defense directly and exclusively to individual actions, but neither does it attempt to extend the theories’ applicability from individual self-defense to national self-defense.

### III.

Joint actions differ from individual actions in that, while the latter are performed by single agents, the former are performed by multiple agents and have individual actions as components. Here is a homely example that illustrates this distinction:

Dale’s car has a dead battery. Roy offers to help with her problem by connecting her battery to his by means of jumper cables, and Dale accepts his offer. When

Roy completes the connection, he signals Dale who is at the controls. She engages the starter and the car starts.

The car is started by Dale and Roy. However, while the proposition that Dale and Roy started the car is true, the proposition that Dale started the car and Roy started the car is false. We might interpret the former proposition as implying that a pair of people started Dale's car. But if we mean by this that the car was started by a set containing Dale and Roy as members, then we would be mistaken. Since sets are abstract entities, they can't be the agents of actions like starting cars. The sentence "Dale and Roy Started Dale's car" means just what it says: there is an action that consists in the starting of Dale's car, and Dale and Roy are its agents.

The idea here can be explained in light of the following, more detailed version of the example:

Dale wants her car to be started. She believes that, if she engages the starter and Roy does his part, then the car will start. Roy wants Dale's car to be started. He believes that, if he connects her battery to his and Dale does her part, then the car will start. Dale's desires and beliefs lead her to engage the starter (call this action A). Roy's desires and beliefs lead him to connect the two batteries (call this action B). Assuming that the car is otherwise in working order and the cables are properly connected, Dale's performing A initiates a sequence of events that merges with the sequence of events initiated by Roy's performing B, forming a sequence of events that results in the starting of Dale's car.

While it isn't true that Dale started her car, and it isn't true that Roy started her car, each of them performs an action (A and B respectively) that is a component of the joint action

that consists in starting Dale's car. They exercise joint agency in doing so, in virtue of the common contents of the beliefs and desires that lead them to perform the individual components of their joint action. Although Dale and Roy's joint action and joint agency are distinct from their individual actions and exercises of agency, the former are explicable in terms of the latter. As a result, the only agents involved in the example are concrete individuals.

What has been said here about the Dale/Roy example can be generalized as follows:

x and y act jointly in bringing about state of affairs S if and only if (a) x and y desire that S obtains, and each believes that there is another person who also desire that S obtains; (b) x and y each believes that, if she acts on these beliefs and desires, and if the other person does so as well, then S will obtain; (c) X's beliefs and desires lead her to perform action v, and y's beliefs and desires lead him to perform w; (d) x's performing v and y's performing w initiate causal sequences that merge to form a sequence that results in S's obtaining.

The joint agency exercised by x and y consists in the common contents of the beliefs and desires that initiate the merging causal sequences that produce S.<sup>12</sup>

This account of joint action and agency can straightforwardly be extended to situations containing more than just two agents.<sup>13</sup> Moreover, joint actions can be composed not only of individual actions, but also of other joint actions. If, in our example, some friend of Roy's helps him connect the cables and does so with appropriate beliefs and desires, then their joint action is a component of a larger joint action performed by Dale, Roy, and the friend.

Because joint actions admit of this sort of structuring, they can be quite complex, and involve large numbers of agents. As was noted above, plays and games provide contexts within which complex joint actions are commonly performed. So do construction projects, sessions of legislative bodies, and battles. Regardless of the composition of a joint action, however, its agents are always concrete individuals exercising joint agency.

Now recall our “Accessory” example, in which Aggressor and Accessory both want Aggressor to kill Victim, and in which he will do so if and only if Accessory isn’t prevented from refueling Aggressor’s tank. If Third Party were to refrain from acting, then the proposed account of joint action and agency would imply that Victim’s death would result from a joint action performed by Accessory and Aggressor. Theirs would be a joint action in virtue of the merging causal sequences resulting from their individual actions, and in virtue of the common contents of the beliefs and desires with which these actions would respectively be performed.

If Third party were to destroy Accessory’s truck before he could refuel Aggressor’s tank, then Third Party’s action would be preemptively other-defensive. While killing Accessory wouldn’t preempt an aggressive action on his part, it would preempt a jointly aggressive action on the part of Accessory and Aggressor.

We can now return to our original example of Al and the terrorist. Let us assume that, if nothing is done to stop him, the terrorist will provide explosive devices to suicide bombers who can’t be prevented from detonating their devices and thereby killing many innocent people. This case is similar in obvious and significant respects to our “Accessory” example. That is, the terrorist bomb-maker and the suicide bombers (and

perhaps others as well) are performing a jointly aggressive action that will result in the deaths of innocents unless something is done to prevent this joint action from being performed.

And there is only one way in which to prevent its performance: Al must kill the terrorist. If Al does so, then his targeted killing is an act of preemptive other-defense against joint aggression. And, in the real world, where targeted killings are planned and carried out by multiple agents, the killing would be jointly defensive.

Having explained how targeted killings can count as defensive homicides, we can now consider whether there are conditions under which they are morally permissible. Doing so will require examining the moral properties of joint actions.

Like individual actions, joint actions can be morally permissible, impermissible, or required. They can also be actions that their agents have a right to perform, or actions that violate the rights of others. Moreover, joint actions possess moral properties in virtue of possessing the same nonmoral properties that determine the possession of moral properties by individual actions. If, for example, a joint action would result in the deaths of innocent people, then it violates the rights to life of those people, and is therefore morally impermissible (*ceteris paribus*).<sup>14</sup>

Additionally, the agents of joint actions can possess moral properties in virtue of the nature of their agency. In particular, they can be morally blameworthy or praiseworthy for contributing to the performance of joint actions which themselves have relevant moral properties. If, say, a joint action is morally impermissible, and if its agents perform their individual actions with appropriately bad intentions and lack excuses for what they do, then they are blameworthy for their contributions to the joint action.<sup>15</sup>

Now, in the realm of individual actions, theories of self-defense typically imply that defensive homicides are permissible only as responses to actions that possess certain sorts of moral defects. According to Thomson, for example, an action possesses the relevant sort of moral defect if it will violate someone's right to life if its agent isn't killed. And according to Susan Uniacke and Jeff McMahan, the relevant defect consists in an action's posing a certain kind of threat to others.<sup>16</sup> The moral defects in individual actions to which these theories refer can also be present in joint actions. Specifically, killing one or more agents of a joint action might be necessary and sufficient to prevent the action from violating someone's right to life. Or a joint action might pose the sorts of threats to which Uniacke's and McMahan's theories refer.

Hence, the theories to which I have alluded could naturally and plausibly be extended from individual defense to joint defense. Rather than attempting to develop any of these possibilities for theories proposed by others, however, I will do so for one that I have defended on a number of occasions.<sup>17</sup>

This theory focuses on situations in which individuals face "closed choices" in the distribution of harm.<sup>18</sup> In a closed choice situation, an individual *x* can't prevent harm from befalling some members of a group *G* (that might include *x*), although *x* can determine which members of *G* are harmed. The theory implies that, if some member of *G* culpably created the closed choice situation, then - *ceteris paribus* - *x* is morally permitted as a matter of justice to distribute the harm to that individual.<sup>19</sup>

Closed-choice situations can clearly be created by multiple agents acting jointly, and these agents can be culpable for doing so. Justice permits the harm to

be distributed to as many of these agents as is necessary to prevent it from being inflicted on innocent potential victims. Suppose, for example, that three individuals jointly and culpably create a situation in which x can prevent the death of an innocent person by – and only by - causing the death of one or more of the individuals who created the situation. Then x is permitted to cause the deaths of as many of these individuals as is necessary to defend their intended victim. This same line of reasoning is applicable to our “Accessory” example. It also applies to our original case of Al and the terrorist, and it implies that Al is morally permitted to kill the terrorist.

The more general implication is that targeted killings are morally permissible if they are relevantly similar to Al’s killing of the terrorist in the circumstances that we are currently envisioning. These are killings that satisfy the following conditions: a number of individuals are culpable for jointly creating a closed-choice situation in which killing some number of them is necessary and sufficient to prevent the loss of innocent lives; as many of the individuals who created the closed-choice situation can be permissibly targeted and killed as is necessary to prevent the loss of innocent lives.

#### IV.

Thus far, no explicit attention has been devoted here to the most difficult and vexing of the moral problems that are associated with targeted killings – namely, that they can result in the deaths of innocent bystanders. The more

general problem, of course, is that of explaining when – if ever – killing (or otherwise causing the deaths of) innocent bystanders is morally permissible. I'll confess up front that my proposals regarding how to solve these problems will be both sketchy and tentative. Nevertheless, the problem of innocent bystanders does deserve some attention in the present context.<sup>20</sup>

We have assumed throughout that AI must kill the terrorist to prevent him and others from jointly causing the deaths of many innocent bystanders. Now let us also assume that, although killing the terrorist will prevent the deaths of these bystanders, it will result in the deaths of others. What we need is a principled basis on which to determine whether AI's killing the terrorist is morally permissible in these circumstances.

According to some writers, causing the deaths of innocent bystanders is never morally permissible. If presented with situations like the one on which we are now focusing – that is, situations in which people face closed choices among the lives of innocent bystanders - these writers are likely to maintain that refrainings aren't causally efficacious. On this view, refraining from preventing someone's death doesn't contribute causally to her death. The refraining is therefore morally innocuous in and of itself; and people confronted with closed choices among innocent lives are morally required to do nothing.

Among writers who believe that refrainings can be causally efficacious, some claim that the moral presumption against causing a bystander's death by acting is stronger than the moral presumption against causing a bystander's death by refraining. This view implies that someone who can prevent the death of one

bystander only by killing (or otherwise actively causing the death of) another, is required to refrain from acting. According to a more restricted version of this position, the moral presumption against killing bystanders is stronger than the moral presumption against preventing others from killing bystanders.<sup>21</sup> This view applies to our example of AI and the terrorist as we are currently construing it. It implies that AI is required to refrain from killing the bystanders whose deaths would be unavoidable if he were to kill the terrorist.

All of these positions imply that targeted killings are morally impermissible if they result in bystanders' deaths. In contrast, both the Doctrine of Double Effect, and the proposition that there exists a moral presumption in favor of minimizing innocent deaths, accommodate the idea that targeted killings are sometimes permissible even if bystanders are killed in the process.<sup>22</sup>

As they apply to cases of the sort on which we are currently focusing, two of the views I have sketched strike me as intuitively more plausible than the others: that there exists a moral presumption in favor of minimizing innocent deaths, and that the moral presumption against killing bystanders is stronger than the presumption against preventing others from killing bystanders. When these two presumptions conflict (as they might very well in targeted-killing situations), it seems to me that intuition points in this direction: refrain from killing bystanders unless the net loss of bystanders' lives would be disproportionately large.

According to this suggestion, AI isn't permitted to kill innocent bystanders along with the terrorist unless his not doing so would result in the net loss of a

disproportionately large number of innocent lives. The concept of proportionality is obviously problematic, but no more so in the present context than in others in which it is commonly claimed to play a role. For example, theories of self-defense almost invariably prohibit defenders from doing disproportionate harm to their attackers.

I will close with some brief comments on another problematic concept that is often appealed to in discussions of defensive homicide – namely, the concept of probability.

Much earlier in this discussion I claimed that probability considerations are relevant not to whether actions are morally permissible, but rather to matters of culpability or blameworthiness. Whether a person is blameworthy for performing some action depends on the intentions with which the person acts, and on whether the beliefs that partly compose these intentions are justified. If the latter beliefs are about matters of fact, then whether they are justified depends on probability considerations, and on whether the agent's estimate of those considerations is reasonable.

Suppose now that Al's destruction of the terrorist's car with one of his drone's missiles results in the net loss of a disproportionately large number of innocent lives. Suppose too that Al launched the missile intending to kill only the terrorist, and that – based on his reasonable estimate of the relevant probabilities – his belief that no bystanders would be killed was justified. Then, in the absence of other relevant considerations, Al isn't blameworthy for causing the bystanders' deaths.

Needless to say, estimating the probability of whether a targeted killing will result in a disproportionately large net loss of innocent life is an enormously complex and difficult business. Making determinations of culpability in cases of morally impermissible targeted killings is therefore correspondingly complex and difficult. However, being committed to making the determinations that are possible is essential to the moral-accountability component of our national policies – especially those (like targeted killings) whose implementation is potentially harmful to innocents.

---

<sup>1</sup> Whether targeted killings are morally permissible under other conditions is not a question that I will address here.

<sup>2</sup> Judith Jarvis Thomson, "Self-Defense and Rights," *The Lindley Lecture* (Lawrence Kansas: University of Kansas Press, 1977), 3.

<sup>3</sup> Thomson presents her account of self-defense in "Self-Defense," *Philosophy and Public Affairs*, 20 (1991), 283-310.

<sup>4</sup> A similar result follows from Suzanne Uniacke's account of self-defense. According to Uniacke, as individuals we possess . . . [the right to life] only so far as we are not an unjust immediate threat to another person's life or proportionate interest. (Suzanne Uniacke, *Permissible Killing* (Cambridge: Cambridge University Press, 1994), 196).

---

Hence, someone who is an "unjust immediate threat" to another's life has no right to life, and can be permissibly killed in either self- or other-defense.

<sup>5</sup> Andrew Altman and Christopher Heath Wellman, "From Humanitarian Intervention to Assassination: Human Rights and Political Violence," *Ethics*, 118 (2008), 253.

<sup>6</sup> I argue for some closely related positions in Phillip Montague, "Blameworthiness, Vice, and the Objectivity of Morals," *Pacific Philosophical Quarterly*, 85 (2004), 68-84.

<sup>7</sup> Analogous issues arise in connection with theories of self-defense according to which killing people is permissible if they are "threats" of certain sorts. The question is whether killing  $x$  at  $t$  is permissible because  $x$  is the right sort of threat at  $t$ , even if - regardless of what is done to  $x$  - he would no longer be a threat at some time later than  $t$ .

---

<sup>8</sup> For an extended discussion of the concept of imminence, see Russell Christopher, "Targeted Killing and the Imminence Requirement."

<sup>9</sup> Although these brief remarks on the nature of preemption suffice for my purposes, much broader purposes are served by the more complete account that is provided by Claire Finkelstein in "Targeted Killing as Preemptive Action."

<sup>10</sup> Thomson's interpretation of what it takes to violate rights is extremely broad, however, as is evident from her claim that "agency is . . . [not] required for violating a right." (Thomson, "Self-Defense" 302.) Arguably, Thomson's account implies that killing x counts as permissible defensive homicide if killing x is the only way to disrupt a causal sequence that would otherwise result in an innocent's death, even if x's role in that sequence is purely passive. Perhaps, then, Thomson would regard her theory as encompassing Accessory's

---

action. After all, Accessory plays an active role in a causal sequences that will result in Victim's death if Accessory isn't himself killed. And similar remarks might apply to the terrorist leader and his potential role in the deaths of innocents.

The matter is, if anything, even less clear when examined in the light of Uniacke's theory. She maintains that defensive force is morally permissible only if used against someone who is "presently" a threat, because "the positive right of self-defense is grounded in the fact that force directly blocks the infliction of unjust harm."

(Uniacke, 185-86) These remarks certainly seem to preclude the possibility of using Uniacke's account as a basis for establishing the permissibility of killing either Accessory or the terrorist leader in our example. But Uniacke also maintains that the use of force against "contingent threats" can be defensive and permissible, and that contingent

---

threats include those who facilitate or assist immediate threats. (Uniacke, 169)

According to Jeff McMahan, if a person is "morally liable to defensive harm," then killing the person doesn't wrong him. He explains the former expression as follows:

the criterion of liability to defensive killing is moral responsibility, through action that lacks objective justification, for a threat of unjust harm to others, where a harm is unjust if it is one to which the victim is not liable and to which she has not consented. Jeff McMahan, "The Basis of Moral Liability to Defensive Killing," *Philosophical Issues*, 15 (2005), 394.

McMahan's position might accommodate the idea that both killing Accessory and killing the terrorist leader are permissible, but determining whether it does would require an explanation of the concept of responsibility for a threat of harm.

---

What has been said here about the implications of these theories for the "Accessory" example, also applies to cases involving accomplices. Here is such a case: Aggressor will shoot and kill Victim using the gun on a tank driven by Accomplice; Aggressor is hidden, but Third Party has a clear shot at Accomplice; if Third Party shoots and kills Accomplice, Victim will have time to escape - and this is the only way that Victim's life can be saved.

Let me hasten to add that my own account of self-defense is no clearer on these issues.

<sup>11</sup> According to Michael Walzer, "states . . . possess rights more or less as individuals do . . . it is possible to imagine a society among them more or less like the society of individuals." (Michael Walzer, *Just and Unjust Wars* (New York: Basic Books, 1977), 58) Since political communities can't act defensively, however, they can't exercise the right of self-defense. It is therefore

---

unreasonable to maintain that political communities have a right of self-defense - or any other exercisable rights.

In a very different context, Peter French attributes actions and agency to "collectivities," and to corporations in particular. (Peter French, *Collective and Corporate Responsibility* (New York: Columbia University Press, 1984,)) Although French never directly confronts the question of whether these entities are abstract, he does claim that they have *members*. If the membership relation - as it pertains to collectivities - differs from the membership relation that pertains to sets, then French needs to explain this difference. In particular, he needs to explain how collectivities could have members without being abstract; and if they could not, then how collectivities could be abstract and still be capable of exercising agency in the performance of actions.

---

<sup>12</sup> Joint agency is actually a bit more complex than these remarks suggest because a person's actions can be initiated by more than one pair of the person's beliefs and desires. For example, Roy's connecting the battery cables might be caused not only by his desire to start Dale's car and his belief that his connecting the cables will help cause this to happen, but also by his desire, say, to check on his memory of how correctly to connect battery cables, together with the corresponding belief.

<sup>13</sup> My explanation of joint actions resembles in certain respects Michael Bratman's account of "shared cooperative activities." ("Shared Cooperative Activity," *Philosophical Review*, 101 (1992), 327-342.) Bratman's account is more restrictive than the explanation of joint activities presented here. For example, the activities of soldiers on one side in a battle could be a joint activity without being a shared

---

cooperative activity. (Bratman actually uses the expression "joint action" in his account of shared cooperative activities, although without explaining what joint actions are.)

<sup>14</sup> The moral properties of a joint action are related to the moral properties of its individual components in complex ways that - fortunately - can be ignored here.

<sup>15</sup> We can say that the agents of the joint action are jointly blameworthy for what they do as long as we don't confuse this notion with that of collective responsibility as the latter is typically understood in the literature. Collective responsibility presupposes collective action - that is, action by collections (or collectivities, to use French's expression), but the only actions presupposed by joint blameworthiness (or praiseworthiness, or responsibility) are those of concrete individuals.

---

<sup>16</sup> McMahan's theory also refers to moral responsibility for creating threats, where moral responsibility is clearly a property of agents rather than of actions. Whether it is a *moral* property, however, is unclear. McMahan emphasizes that moral responsibility isn't culpability or blameworthiness, and so he might simply be identifying a kind of responsibility that differs from mere causal responsibility. If so, then being morally responsible for posing a threat might be equivalent to posing it freely, or voluntarily, or intentionally - implying that moral responsibility isn't a moral property.

<sup>17</sup> This theory is developed at length in "Self-Defense and Choosing Among Lives," (*Philosophical Studies*, 40 (1981), 207-19) and in *Punishment as Societal Defense* (Lanham, Maryland: Rowman and Littlefield, 1995).

---

<sup>18</sup> I originally referred to these as "forced choices," but this terminology led to misinterpretations of my position.

<sup>19</sup> The *ceteris paribus* conditions referred to here pertain to proportionality, to doing the minimum harm necessary to accomplish the permitted or required distributions, and to "side effects" such as harm to innocent bystanders.

Although I have formulated the theory in terms of reference to moral permissibility, a more precise formulation would refer to moral rights and moral requirements. That is, in culpably created closed-choice situation, people have moral rights to favor themselves, and are morally required to favor other innocents. In virtue of its reference to culpability, this theory implies that defensive homicide - understood as a moral right or a moral requirement - is a response not only to morally defective actions, but also to morally defective agency. In closed-choice situations that are

---

impermissibly but not culpably created, defensive homicide is merely permissible. (This position is defended in Phillip Montague, "Self-defense and Innocence: Aggressors and Active Threats," *Utilitas*, 12 (2000), 62-78.)

<sup>20</sup> The question of how to explain the concept of an innocent bystander won't be addressed here.

Present purposes are served by noting that young children are always innocent bystanders, so the problem of concern here can be interpreted as that of determining when - if ever - causing the deaths of young children is morally permissible.

<sup>21</sup> A principle similar to this is central to Bernard Williams's well-known critique of utilitarianism. (J.J.C. Smart and Bernard Williams, *Utilitarianism for and Against* (Cambridge: Cambridge University Press,, 1973).

<sup>22</sup> I have stated the second of these positions in terms of a presumption in favor of minimizing innocent deaths rather than a requirement to do so,

---

because the latter would be vulnerable to familiar counterexamples.