

Spring 1-1-2015

Drones, Intention, and Psychological Harm: An Argument for Restraint

Tucker Allen Marks

University of Colorado Boulder, tucker.marks@colorado.edu

Follow this and additional works at: https://scholar.colorado.edu/phil_gradetds

 Part of the [Ethics and Political Philosophy Commons](#), and the [Philosophy of Science Commons](#)

Recommended Citation

Marks, Tucker Allen, "Drones, Intention, and Psychological Harm: An Argument for Restraint" (2015). *Philosophy Graduate Theses & Dissertations*. 46.

https://scholar.colorado.edu/phil_gradetds/46

This Thesis is brought to you for free and open access by Philosophy at CU Scholar. It has been accepted for inclusion in Philosophy Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

DRONES, INTENTION, AND PSYCHOLOGICAL HARM:

AN ARGUMENT FOR RESTRAINT

by

TUCKER ALLEN MARKS

B.A. Colgate University, 2013

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirement for the degree of

Master of Arts

Department of Philosophy

2015

This thesis entitled:
Drones, Intention, and Psychological Harm: An Argument for Restraint
written by Tucker Allen Marks
has been approved for the Department of Philosophy

Alastair Norcross

Chris Heathwood

Date 05/29/15

The final copy of this thesis has been examined by the signatories, and we
Find that both the content and the form meet acceptable presentation standards
Of scholarly work in the above mentioned discipline.

Marks, Tucker Allen (M.A., Philosophy)

Drones, Intention, and Psychological Harm: An Argument for Restraint

Thesis directed by Associate Professor Alastair Norcross

The use of uninhabited aerial vehicles, or drones, by the United States military has drawn increasing attention in recent years, and with that attention has come broad international opposition. Despite this opposition, it can be difficult to figure out what, if anything, is wrong with the use of drones. They can reach areas that would not be accessible to human soldiers, they can survey a location for hours and days at a time, and they can carry out strikes that would be too dangerous to carry out through other means. Further, they are thought to be more precise than other weapons, and therefore better able to avoid civilian casualties. One element that is a focus of popular criticism of drones, yet has failed to appear in philosophical literature, is the psychological harm drones can cause. Drones have been shown to cause significant amounts of fear in civilian populations, which should be considered when discussing the permissibility of their use. Here, I argue that this psychological harm makes the use of drones impermissible in many cases. Relevant to this conclusion is the role of intention: frequently, the fact that harm is unintended has been used to deny its relevance to permissibility. I consider several theories of civilian targeting, and argue that we should accept an account based in liability and/or desert rather than one that depends on intention. If we accept this sort of view, we seem to get more intuitive and more consistent results in a variety of cases, including cases involving drones. We are better able to explain what is both distinctive about drones and what makes their use impermissible in many cases.

CONTENTS

CHAPTER

I.	INTRODUCTION: AN ARGUMENT FROM PSYCHOLOGICAL HARM... 1	
II.	THE DOCTRINE OF DOUBLE EFFECT: HARM, CHARACTER, AND CIVILIAN IMMUNITY.....	12
III.	DOUBLE INTENTION AND SUPREME EMERGENCY.....	22
IV.	LIABILITY, INTENTION, AND HARM.....	27
V.	CONCLUSION: THE PROPER ROLE OF DRONES.....	42
WORKS CITED.....		45

I. Introduction: an Argument from Psychological Harm

The use of unmanned drones¹ in counterinsurgency operations has received significant attention recently, at least in part due to the extensive use of drones by the United States in Pakistan, Yemen, and other Middle Eastern countries. These actions have been controversial, to say the least. International opposition to the use of drones by the United States is strong, with most countries surveyed showing higher percentages opposing our drone use than supporting it (“Widespread Opposition to Drones,” 2014). As drones further become the centerpiece of United States counterterrorist operations, determining the proper policy governing their use gains greater practical importance. Though there are many features of the drone policy of the United States that have been criticized, including civilian casualties, the violation of sovereignty of countries in the region, and our use of “signature strikes²,” I will not be directly concerned with any of these here. These are all features that can be in principle addressed: drone operators might strike only in cases where civilian casualties meet proper proportionality constraints, the United States government might fully seek permission from the countries in which it operates, and signature strikes might be forbidden. The psychological harm caused by drones, however, may be something that cannot be easily avoided; that is, any use of drones may cause psychological harm on a wider group than is physically harmed in a strike. It is this psychological harm that I will investigate here. By psychological burden, or psychological harm, I mean fear and distress suffered by non-combatants living in areas affected by drone strikes. In this section, I will discuss several features of this psychological harm, argue that drones may be more likely to cause these harms than other forms of military action, and, finally, argue that these harms give us reason to place restrictions on the permissible use of drones. The rest of the paper will consist in

¹ Also referred to as unmanned/uninhabited aerial vehicles (UAVs). I will use “drones” throughout.

² This refers to the targeting of people based on “patterns of life,” as opposed to specific knowledge of their identities (“Living Under Drones,” 2012: 12).

a discussion of targeting principles and the role of intention, with the aim of defending this conclusion.

The most prominent defense of drones use has been made by Bradley Jay Strawser, in “Moral Predators: the Duty to Employ Uninhabited Aerial Vehicles.” There, he argues for the truth of the following principle: “For any just action taken by a given military, if it is possible for the military to use UAV platforms in place of inhabited aerial vehicles without a significant loss of capability, then that military has an ethical obligation to do so.” (Strawser 346) In his discussion, he addresses several objections to the use of drones, but does not discuss the psychological harm that drones may cause. He does, however, include a *ceteris paribus* clause in his principle, which would be affected by the arguments here; he notes:

“My argument rests on the premise that if an agent is pursuing a morally justified yet inherently risky action, then there is a moral imperative to protect this agent if it possible to do so, unless there exists a countervailing good that outweighs the protection of the agent.” (Strawser, 2010: 343)

The countervailing good, in this case, would be in avoiding causing this psychological harm. Strawser’s discussion on civilians is focused on the ability for drones to limit physical casualties; he concludes that drones may better limit these casualties than traditional aircraft, and his conclusion may be supported by recent data³ (Strawser, 2010: 351-2; Plaw, 2013: 152). Indeed, Avery Plaw notes that drones “compare favorably with military alternatives in the region” with respect to civilian deaths (Plaw, 2013: 152). The issue of psychological harm to civilians, however, is largely without consideration in philosophical literature on drones. This is

³ It is worth noting that civilian casualty numbers from drone attacks are far from certain, and that many sources have reliability concerns. See Chapter Two of the Stanford/NYU *Living Under Drones* report for a full discussion. I will assume here that drones are reasonably effective in avoiding civilian casualties for the sake of argument. If it turns out that they are not, then we have even greater reason to avoid their use.

particularly surprising due to its central role in the public narrative on drones; the fear of civilians living in areas surveyed by drones is the central focus of the Stanford/NYU *Living Under Drones* report, and has been featured in many public discussions, including a segment on *Last Week Tonight with John Oliver*.

Psychological harm, as a consequence of war, seems unavoidable. In any area in which war is fought there is a risk of civilian casualties, and this risk alone can clearly lead to significant fear and distress. Given that traditional theories of just war and international law all allow for some amount of collateral damage, and that many (perhaps all) armies violate these restrictions in war, this fear seems justified. Though certain forms of attack may lead to more fear among the civilian population than others, it seems that all may cause it in some way. If civilians knew that any given attacking force was committed to limiting civilian casualties, and had precise enough weapons in order to do so, this should certainly lessen the distress they face, but certainly cannot eliminate it (as we shall see).

One factor relevant to this discussion, given international opposition to drone use, that seems to have a clear impact on the potential for a military action to inflict psychological harm is the perceived result of similar actions in the past. If a military consistently kills civilians in its operations, for instance, any civilian population would rightly fear that military's arrival. The same would hold for a military (that otherwise avoids civilian casualties) using a weapon that is widely thought to have a history of increased harm to civilians. Critically, while the actual level of civilian casualties is likely to lead to this harm, it is the perceived level of casualties that has a far greater impact: even if the actual history of a particular sort of action contains very few civilian casualties, factors from increased media attention to prominent mistakes can create a very different impression upon the civilian population in an area. A weapon that is actually safer

than its alternatives may cause a greater level of fear and distress, making it more psychologically harmful.

Drones themselves seem clearly to cause a high level of psychological harm. A report by the International Human Rights and Conflict Resolution Clinic of Stanford Law School and the Global Justice Clinic at New York University School of Law, “Living Under Drones: Death, Injury and Trauma to Civilians from US Drone Practices in Pakistan,” offers the following, rather horrifying, picture:

“Importantly, those interviewed for this report also described how the presence of drones and capacity of the US to strike anywhere at any time led to constant and severe fear, anxiety, and stress, especially when taken together with the inability of those on the ground to ensure their own safety. Further, those interviewed stated that the fear of strikes undermines people’s sense of safety to such an extent that it has at times affected their willingness to engage in a wide variety of activities, including social gatherings, educational and economic opportunities, funerals, and that fear has also undermined general community trust. In addition, the US practice of striking one area multiple times, and its record of killing first responders, makes both community members and humanitarian workers afraid to assist injured victims.” (“Living Under Drones,” 2012: 55)

While many weapons used in war cause psychological harm, this description is largely unique to drones. While evidence surveyed by Avery Plaw indicates that drones are no less precise than other forms of attack (and may indeed be better at avoiding civilian casualties), drones fire explosive ordinance that can give the impression of imprecision (Plaw, 2013: 126). The following account should help illustrate:

“Interviewees described the experience of living under constant surveillance as harrowing. In the words of one interviewee: ‘God knows whether they’ll strike us again or not. But they’re always surveying us, they’re always over us, and you never know when they’re going to strike and attack.’” (“Living Under Drones,” 2012: 81)

This account focuses on the ability for drones to strike at any moment, as do many others in the report. This seems to highlight the biggest difference in ability between drones and manned aircraft: while manned aircraft generally cannot hover over an area for an extended period of time, due to the presence of a human pilot, drones can remain over an area for up to twenty-four hours, giving the impression of constant surveillance (Strawser, 2010: 355; “Living Under Drones,” 2012: 9). While soldiers on the ground can remain in an area for an extended period, they generally are not using missiles to eliminate their targets, and their presence is not constantly announced by a noise overhead.

Highlighting the apparent inescapability of drones, David Rohde, a journalist kidnapped by combatants in Pakistan, reported:

“From the ground, it is impossible to determine who or what they are tracking as they circle overhead. The buzz of a distant propeller is a constant reminder of imminent death. Drones fire missiles that travel faster than the speed of sound. A drone's victim never hears the missile that kills him.” (Rohde, 2012 , also cited in “Living Under Drones,” 2012: 80)

The features Rohde describes all help drones to be more effective in fighting enemy combatants, but clearly contribute to the psychological harm of civilians living in the area as well. While there are other forms of attack that provide similar worries, including, but not limited to, artillery strikes and manned air strikes, those each have technical limitations that prevent them from

causing these same levels of psychological harm. Artillery is not visible or audible overhead, and so cannot create this fear before a strike, and the limitations of manned air strikes have already been mentioned. A further feature of these reports, as well as others, is that they contribute to the negative perception of drones. Even if civilian casualties from drone strikes are in fact rare, reports of civilian casualties, both physical and psychological, will contribute to the fear caused by drones.

While it seems clear that drones can and perhaps do cause greater levels of psychological harm than other forms of military action, it is not immediately clear how this harm is relevant to the permissibility of drone use. Drones reduce the risk of physical harm to our soldiers, and it is not clear how physical and psychological harms should be compared. In what follows, I will discuss the role that the aforementioned psychological harm should play in our assessment of the permissibility of drone use.

An important distinction here is between two sorts of military action: those that occur in the course of “standard” war, and those that are aimed at non-state actors in a counterinsurgency setting. The psychological harm described above is likely to be prevalent in standard war regardless of weapons used: if a village is caught in combat between two armies, drones will likely be among the least of civilian worries. Artillery strikes, with the potential to be far less precise than drones, become a significant danger, and there is the risk of large-scale fighting taking place among civilian residences; these are all likely to have an equal or greater propensity to disrupt the lives of civilians. While drones may well add to this fear for the reasons mentioned above, it is difficult to accept that this addition would be significant. Drones deployed away from the front may cause harm similar to that in counterinsurgency cases; due to their location away from the front, however, alternative strategies will be more limited.

The use of drones in counterinsurgency operations is a more difficult issue. Unlike cases of open war between countries, where psychological harm may be unavoidable, it seems quite avoidable in these cases of more limited military action. Each of the accounts offered above is in reference to drones being used to attack non-state combatants, not soldiers in a state army. The behavior of the non-state combatants toward the local population will likely determine the “default” level of fear in a particular area; if the above accounts from “Living Under Drones” are any indication, the presence of drones does significantly change the status quo experienced by civilians in the relevant areas of Pakistan. Unlike in the case of conventional war, it seems likely that other methods of attacking insurgents may generate less psychological harm in civilians without increasing the risk of physical harm; e.g. deploying a small team of soldiers with particular training to avoid civilian casualties. Given the status of psychological harm as some form of collateral damage, it might seem that we are obligated to choose these other methods. Whether we are in fact obligated to choose those other methods will depend on which theory we accept regarding civilian harm. The primary candidates are the Doctrine of Double Effect (with various modifications), Act Utilitarianism, and liability theories⁴.

While the bulk of this paper is devoted to discussing these theories, I will quickly discuss Act Utilitarianism here. Its assessment of drone use is quite straightforward: if drone use causes more harm than alternative means of achieving the same outcome in a particular case, then it is impermissible in that case. I have suggested here that the psychological harm caused by drones might cause it to fail this condition; however, whether or not this is the case is an empirical matter. If, for example, claims that drones are less likely to cause civilian deaths are true, then there may well be cases where the decreased chance of civilian death outweighs the

⁴ I include under “liability theories” both deontological theories dealing with liability and consequentialist theories that include adjustments for desert.

psychological harm. Given that actual civilian deaths (in cases involving drones and other strategies) seem to be controllable to a greater degree than psychological harm resulting from drone use, it seems that we may well reach a point where drones are nearly guaranteed to cause more harm than at least some alternatives, though this will obviously vary between cases. In some cases, of course, drones will be the only option, and so the benefits of the operation will determine its permissibility; for instance the combatants may be inaccessible without drones. But in many cases alternatives will be feasible. While I will discuss other theories in more detail later, the general claim that one ought to reduce collateral damage when possible is endorsed by all of them, and these considerations support the conclusion that we can indeed reduce collateral damage by not using drones, and so we ought not to use drones.

This conclusion, however, fails to account for the increased risk to soldiers engaging in counterinsurgency operations. Importantly, any operation where we choose not to use a drone is one where at least one person in our military is placed at increased risk of physical and/or psychological harm; after all, at least one person will be present in the area who otherwise would not have been. It is the ability to decrease risk to our soldiers that motivates Strawser's argument in favor of the use of drones. Strawser defines the Principle of Unnecessary Risk as follows:

“If X gives Y an order to accomplish good goal G, then X has an obligation, other things being equal, to choose a means to accomplish G that does not violate the demands of justice, make the world worse, or expose Y to potentially lethal risk unless incurring such risk aids in the accomplishment of G in some way that cannot be gained via less risky means.” (Strawser, 2010: 344)

This principle should, on Strawser's account, oblige us to use drones, as their existence makes the risk to our soldiers unnecessary in many cases. As Strawser admits, however, this all relies

on the *ceteris paribus* clause in the Principle of Unnecessary Risk: in many of these cases, other things will not be equal due to the psychological harm caused by drones. Several philosophers have endorsed the claim that, in at least some cases, soldiers must take the risk of additional harm on themselves in order to protect civilians; I will discuss some of these claims later, but for the moment will add that this result is intuitive. If, for example, a civilian population was fully innocent, that is to say not at all aiding the unjust combatants (and not responsible for the conflict), and the unjust combatants are harming them by occupying their village, the claim that just soldiers should avoid causing even further harm to these civilians does not seem unreasonable, even if doing so requires the just soldiers to take on additional risks. Obviously, this issue is not settled, but I will return to it later in the paper.

This result, if true, may still appear strange; after all, the potential harm to the soldiers is physical, while the harm to civilians is psychological. This difference is important: it makes justifying the increased risk placed on the soldiers more difficult. Given two factors noted earlier in this paper, however, it seems we do have this justification in some cases. First, while the harm to civilians is psychological, it may still be significant; while balancing different kinds of harm is difficult, there is no reason to think that psychological harm could not be as morally significant as physical harm, even death, if the psychological harm is severe enough and inflicted upon a large enough number of people. Drones seem especially prone to inflicting this psychological harm through their very presence. Second, the position of civilians who lack responsibility for the actions of the insurgents is a particularly unenviable one; even if they are not actively being threatened by the insurgents, they know that attacks on these insurgents may easily result in collateral damage. Finally, the potential physical harm to the soldiers is just that: potential harm. The psychological harm caused by drones is, as I have attempted to indicate here, effectively

certain. It is dependent simply on there being civilians in the broad area, not on the particular aiming of the drone operator. In most cases, there will merely be a risk of harm to the soldiers that would have to be deployed instead of a drone. This risk should be taken seriously, but the fact that it is a risk should cause it to weigh less in the comparison.

The following principle seems to be a reasonable conclusion from the above discussion: in cases where the use of a drone would cause civilians greater harm, physical or psychological than an alternative form of action, the civilians bear little responsibility for the actions of the target insurgents, and the added risk to our soldiers from choosing the alternative action is not excessive, we ought not use drones. Conversely, we are permitted to use drones, all else equal, when either 1) there is not an increase in harm from their use (this covers relevant applications in conventional war); 2) the civilians in the area are voluntarily providing support to the insurgents; 3) The increased risk to our soldiers from the alternative action significantly outweighs the added psychological harm from using the drone. Each of these assumes that the physical risk to the civilian population is at least roughly equivalent between the use of a drone and the alternative action, whatever that happens to be; if the use of a drone poses significantly lower physical risk to both civilians and our soldiers, that is strong reason to prefer it over the alternative. Following this principle may serve at least two purposes: morally, it accounts for an element of civilian harm that does not receive significant philosophical attention, and prudentially, it may help create a positive image of our counterinsurgent actions. Given the inherent potential for drones to cause psychological harm, as well as the harm they appear to have actually caused, we should restrict our use of drones in the cases identified by this principle. I will now discuss the assessment of drone use by the Doctrine of Double Effect and the plausibility of that doctrine;

my conclusion here depends on the permissibility of harming civilians in certain ways, and so it is critical that we find the right principle to assess these harms.

II. The Doctrine of Double Effect: Harm, Character, and Civilian Immunity

In the above section I have, I hope, shown that the use of drones involves a far greater amount of harm to civilians than has previously been accepted. If I am to move from that conclusion to the conclusion that we ought not to use drones in many cases, I must show that the best targeting theory reaches such a conclusion. Here, I will be discussing perhaps the most common targeting theory in the just war tradition: the Doctrine of Double Effect. First, I will detail the theory as it is generally given, and use it to assess the use of drones. Following this, I will consider reasons for accepting and rejecting this doctrine, and ultimately conclude that we ought to reject it as the proper account of targeting in war.

The Doctrine of Double Effect permits an action with bad effects if that action meets the following four conditions, as stated by Michael Walzer:

- “1) The act is good in itself or at least indifferent, which means, for our purposes, that it is a legitimate act of war.
- 2) The direct effect is morally acceptable – the destruction of military supplies, for example, or the killing of enemy soldiers.
- 3) The intention of the actor is good, that is, he aims only at the acceptable effect; the evil effect is not one of his ends, nor is it a means to his ends.
- 4) The good effect is sufficiently good to compensate for allowing the evil effect; it must be justifiable under Sidgwick’s proportionality rule.” (Walzer, 1977: 153)

The critical condition is the third, as it summarizes the practical use of the doctrine: an action is permitted to have a bad effect so long as that effect is unintended as an end or a means. The bad effect may be foreseen, so long as it is not intended. Walzer is unclear in his description between two meanings: this condition can be understood as a subjective restraint on intending the evil

effect as a means, or as an objective condition on the evil effect actually being a means. I will understand it as the former; that is how it is generally meant, as it is a constraint on intentions. It is this condition that separates the Doctrine of Double Effect from Utilitarianism (embodied in the fourth condition) or an absolute prohibition on bad effects. The relevant bad effect for our purposes is harm to civilians. Harm to soldiers, while clearly bad insofar as it is harm, is considered a “good effect” in that it prevents the harm done by those soldiers.⁵ Civilians, by comparison, are taken to pose no threat. There is one further feature that is, I think, entailed by the proportionality condition: the unintended bad effect must be as small as is consistent with achieving the intended end; put differently, the bad effect must be necessary. The bad side effect can only be said to be proportional if it leads to the greatest amount of net good possible, which means minimizing any unintentional harms.

How, then, should the use of drones be judged by the Doctrine of Double Effect? It may well be permissible. The intended targets of drones are (in paradigm cases) members of opposing militaries or terrorist groups, while the harm, physical and psychological, done to civilians is merely foreseen, but not intended. So long as the good done by attacking the intended target outweighs the harm to civilians, it seems clear that drone use passes the requirements of the Doctrine of Double Effect. The addition of psychological harm to this situation, however, does raise questions about this conclusion. As I discuss above, the fact that drones cause less overall harm to civilians than other options is less likely, given that the estimated harm from drone use has increased, and if it turned out that other options did in fact cause less harm to civilians, then drone use would be impermissible. This is, of course, an empirical question, and one that cannot be settled here, but I think it is worth noting that it is at least less likely that the Doctrine of

⁵ This is a departure from Aquinas’ original statement of the doctrine; there, even the death of an assailant killed in self-defense is considered a bad effect that must not be intended. See Aquinas, *Summa Theologica* II.II, 64, 7.

Double Effect would approve of drone use once psychological harm done to civilians is considered.

I am happy to grant, for the sake of this discussion, that drone use would be permitted by the Doctrine of Double Effect. The more important question seems to be this: should we rely on the Doctrine of Double Effect to determine the permissibility of actions in war? If we are to answer the question “should we be using drones?” we must answer this first.

What follows is the standard case used to motivate the Doctrine of Double Effect. Suppose country A is fighting a just war against country B. A has discovered a munitions depot of B, and determines that the destruction of the depot will aid significantly in their effort to win the war. Located next to this munitions depot is a school, which will be destroyed in any bombing of the munitions depot. A must decide between two missions to run against this depot: the first, called “tactical bomber,” will target the munitions depot in an attempt to reduce B’s supplies, and thus help A win the war; as an unintended effect, the school will also be destroyed, killing, suppose, 100 children and adult civilians. The alternative, called “terror bomber,” involves the exact same action, namely the bombing of the munitions depot. But here, the intended effect is demoralizing the enemy through the destruction of the school and the civilians within, and thus aiding A’s war effort, while the destruction of the munitions depot is merely an unintended side effect.

While the tactical bomber and terror bomber take the same action, namely the bombing of a particular location within country B, each has a different intention in doing so. The common intuition in this case is that the tactical bomber acts permissibly, while the terror bomber does not. This is the very same conclusion reached by the Doctrine of Double Effect. Because the actions are otherwise identical, it must be the case that the difference in intention results in any

difference in permissibility. This is, effectively, the central argument⁶ in favor of the Doctrine of Double Effect's third condition: we have a strong intuition that cases like tactical bomber are permissible while cases of terror bomber are impermissible, and the only principle that gives that result is one that appeals to the relevance of intention, which is precisely the point of this condition of the Doctrine of Double Effect. The only way to avoid the argument, it seems, is to reject the intuition that there is a difference in permissibility in the tactical bomber and terror bomber cases.

While I admit that I do have this intuition, I am willing to reject it for reasons I will soon discuss. First, however, I would like to consider a more forceful variant of this argument offered by Jeff McMahan in his article, "Intention, Permissibility, Terrorism, and War." McMahan makes a more directly practical claim than above: not only is something like the third condition of the Doctrine of Double Effect required to distinguish between tactical bomber and terror bomber cases, but it is required to distinguish between just war and terrorism in general⁷.

McMahan thinks making this distinction is of great importance, as he states:

"Intuitions about hypothetical trolley cases are, of course, important for moral theory, but our beliefs about the morality of war and terrorism affect whether and how we pursue these activities, and this is obviously a matter of the greatest practical significance.... But if intention does not magnify the moral objection to killing an innocent person — if, that is, an innocent person's right not to be killed imposes no stronger constraint against intentional killing than it does against foreseen but unintended killing — then terrorism

⁶ Different trolley cases are also used in this argument, which I will discuss shortly.

⁷ McMahan understands terrorism to be "the intentional harming (usually killing) of innocent people as a means of intimidating and coercing other people associated with them, usually for political purposes." (McMahan, 2009: 360) This definition is close to the standard one, though it is worth noting that it is not universally accepted. For instance, many consider attacks on US military locations to be terrorist attacks, even when there is clearly a military target. Though I think we need a better definition of terrorism, I will accept McMahan's here.

should be no more objectionable, other things being equal, than military action in war that foreseeably but unintentionally kills innocent people.” (McMahan, 2009: 359-360)

The argument, then, seems to be that since terrorism is clearly impermissible, yet at least some actions in war are permissible, we must be able to draw a clear distinction between them. The Doctrine of Double Effect, and views like it, point to one clear way in which terrorism seems different from permissible actions in war, and since no alternative view can capture the distinction, we must accept something like the Doctrine of Double Effect. I agree with McMahan that it is of great practical importance to determine what exactly is wrong about terrorism, and that we must distinguish acts of terrorism from acts of just warfare. But I think we can distinguish between them without requiring a connection between intention and permissibility. Though McMahan considers four possible alternative explanations of the difference between terrorism and just war, I think there are a few other considerations that may still succeed.

If we consider the most infamous instances of terrorism, a couple of common themes emerge. The first is that the sort of killing involved seems, in general, completely unnecessary for accomplishing the (occasionally stated) goal. If we take the most charitable view possible toward groups like Al Qaeda, and accept their claim that they intend to stop the United States from killing innocent civilians in the Middle East, their actions still do not clearly seem to work toward that goal: their attacks are not focused on our ability to fight, nor are they effective in dissuading us. Second, it seems far more plausible that there is some less noble end⁸ toward which most terrorist groups act: it is the badness of the end, along with the perception (true or

⁸ This reads like “intended end,” which it cannot mean if I am arguing against the relevance of intention. Instead, I mean roughly that the overall effect is bad; it would be the difference between the terror bomber doing bad to do good (net positive) and a terror bomber doing bad to do more bad (net negative). If the only means of defeating Nazi Germany was a “terrorist attack,” I doubt it would be criticized. (That terror bombings in WWII are criticized does not tell against this: the main point of criticism seems to be that we made those bombings unnecessarily, rather than that we made them at all.)

not) that the actions of terrorists are excessive (or simply unhelpful) in achieving the end that separates them from just war. A couple of cases may help motivate this. First, consider a world, where one of the United States' top generals was visiting the World Trade Center on September 11, 2001, and the attack on the towers was planned specifically to ensure that the general was killed. It seems doubtful to me that the reaction to the attack would be any different than it was in reality. The same civilians would have died, the same amount of destruction would have been caused, and so on. We can make a similar substitution into one of the cases McMahan discusses, the US bombing of Tokyo during World War II. It would be shocking if our assessment of that bombing changed simply because military targets were placed among the casualties, and declared the intended targets. In each of these cases, it seems that the deaths of civilians, and destruction of non-military areas, are what explains our reaction, rather than the intentions of the attacker.

Another pair of cases illustrates the relevance of the goodness and badness of ends. First, consider a common variant of the trolley case. A trolley is moving down a track toward five people. You are standing on a bridge above the track, and another person is standing next to you. You determine that pushing this person in front of the track would kill her, but would also derail the trolley, and prevent it from killing the five people. According to the Doctrine of Double Effect, it would be impermissible to push the person in front of the trolley, because doing so would be intentionally using her as a means⁹.

⁹ I am assuming (for the sake of keeping the examples simple) that this actually does involve using the person's death as a means to saving the five on the track. One might say "really, it's simply the person's presence that is the means to the good effect, and her death is a foreseen but unintended side effect;" I will leave aside any problem this sort of move might reveal about the Doctrine of Double Effect, and note that we could construct a more complicated case that clearly does involve using the death of the person as a means to a good effect. Such an example would be significantly more complicated, however, and the exact judgment in this case (that pushing the person is impermissible) is not itself critical to my argument.

Second, a structurally similar case can be made in the context of war¹⁰. Suppose country A is fighting a just war against country B. Unfortunately for A, B possesses nuclear weapons, and the general commanding B's military has decided to use them to destroy several of A's cities, against the general wishes of his fellow citizens. A discovers this plot, and (let's suppose) has two options: (1) A can bomb B's capital, killing the general along with 100,000 of B's citizens; or (2) an agent of A can push a civilian in front of the general's car on his way to wherever he must go to launch the nuclear missile, killing the civilian and the general. In either case, B does not launch the nuclear missiles, and 200,000 lives are saved in A. If we follow the Doctrine of Double Effect, we must choose (1), despite this option leading to the unnecessary death of 99,999 additional people¹¹. (1) meets the proportionality constraint, as it saves more people than it kills, and the deaths of the 100,000 in B are all unintended. In (2), the 1 civilian death was the means to killing the general, and so the Doctrine of Double Effect must rule it impermissible. It would be shocking to me, however, if anyone actually preferred (1) to (2), and the preference for (2) despite its violation of the Doctrine of Double Effect can be explained by the good end being pursued: saving 200,000 innocent people is clearly an end worth pursuing, and this can explain why we might be comfortable endorsing (2) here, but would judge a similar action pursuing a bad end as impermissible. Conveniently, this explanation also makes sense of any differing intuitions between this case and the trolley case above: saving 200,000 is much better than saving 5 for the same cost.

Though the discussion so far (hopefully) gives us good reason to deny the Doctrine of Double Effect, it does nothing to explain what is going on in the case of the tactical bomber and terror bomber. In that case, the consequences are, by any measure, equal, and the means chosen

¹⁰ Norcross constructs a functionally identical case in a draft of "The Road to Hell." (Norcross, 2011: 15)

¹¹ Cases like this are used to motivate Michael Walzer's "supreme emergency" doctrine; I will be discussing this doctrine later, but it should suffice to say that a less dramatic case would face the same problem.

are taken to be equally effective in accomplishing their respective ends. The *only* difference between the tactical bomber and terror bomber is intention, and if we wish to deny the Doctrine of Double Effect we must have an explanation as to why people generally take the tactical bomber's action to be permissible and the terror bomber's impermissible that is consistent with the actions' being, in fact, equally permissible. Alastair Norcross, in a draft of his unpublished paper, "The Road to Hell," suggests that we should use intentions to assess character instead of action (Norcross, 2011: 22). The natural move in the tactical/terror bomber case, then, is to say that the willingness to target civilians of the terror bomber is evidence of her worse character, and that this explains our intuition in favor of the tactical bomber. Though the acts are equally permissible, the character of the tactical bomber is praiseworthy in comparison to the terror bomber, and we were simply mistaken in thinking that our intuitions tracked permissibility in that case¹².

Norcross, however, argues that our intuitions are wrong when considered in this way as well, as the terror bomber's character is not clearly worse than that of the tactical bomber. He argues:

"Given certain choices between two different war-ending raids, the strategic bomber would perform the worse action, but the terror bomber would perform the better one.

Furthermore, if the only way to end the war were to carry out the terror raid, the strategic bomber would refuse, thus condemning a million people to death...." (Norcross, 2011:

25)

¹² Or, a virtue ethicist might think that the judgment of praiseworthiness entails a judgment of permissibility. This would not reinforce the Doctrine of Double Effect, as it does not clearly apply to all cases equally: in the above case involving either the intentional killing of one civilian or the unintentional killing of 100,000, it seems obvious that choosing the former indicates a better character than choosing the latter.

I think it is certainly true that the tactical/terror bomber case can easily be translated into a case where the tactical bomber is forced (by character) to choose a worse action, where the terror bomber would (again by character) always choose the best outcome. This is likely the correct analysis of their respective characters.

There are two other very nearby cases, however, that might provide a psychological explanation for our intuitions in the original bomber case. Instead of the original case, where we compare two pilots flying qualitatively identical missions (with different intentions), consider a case where a pilot is presented with a choice between the two missions. She is told all of the mechanics of each mission, the equal number of civilian casualties, the intended goal of each mission (to destroy a munitions depot or to kill the civilians), and the equal effectiveness of each mission in pursuing that goal. She is not told, however, that the missions are identical *except* for the intended goal (suppose she is not given the exact location of the target, for instance). Further, she believes all of the estimates about civilian casualties and effectiveness, so the only remaining choice is between intended ends. Something seems blameworthy, in that case, about choosing the terror bombing. The choice involved seems relevantly similar to the following choice, between two (genuinely) different missions, one where no civilians will be killed, with the goal of destroying a munitions depot; and one with the goal of killing civilians¹³. If the missions are believed to be equally effective, and the pilot chooses to kill the civilians rather than bomb the munitions factory, she clearly makes the worse choice: she chooses to cause harm when doing so has no corresponding benefit. This case is, obviously, different from the original case in that respect, but it does not seem like much of a leap to say that that's what we think about when we judge the character of the tactical bomber and terror bomber. We are tracking the difference in intended end, rather than the actual difference between the cases.

¹³ This is the tactical/terror choice without any collateral damage.

I think we can correct for the difference in overall harm and maintain the same result. If we replace the munitions factory with, say, a military fortress holding the same number of (active duty, ready-for-combat) soldiers as the civilians in the terror case, we equalize the (apparent¹⁴) harm done by choosing either mission. In this case, bomber is choosing between killing soldiers and killing civilians, and I think most people would criticize the bomber for choosing to kill civilians. This could be because they believe the person who chooses to kill the civilians in this case would choose to kill civilians in other, unjustified, cases, or it could be because it is, for reasons unrelated to the Doctrine of Double Effect, impermissible to choose to kill the civilians. After exploring some modifications to the Doctrine of Double Effect, I will defend this sort of view.

¹⁴ Later, I will discuss a view by which the harm done would not actually be equal.

III. Double Intention and Supreme Emergency

Faced with objections to the Doctrine of Double Effect (and similar theories) along the lines that I have pressed, Michael Walzer has offered two principles: the doctrine of supreme emergency, and the doctrine of double intention. I will begin with double intention. Walzer offers a modification to the Doctrine of Double Effect designed to prevent one group of counterintuitive cases, namely those in which the Doctrine of Double Effect seems to permit excessive civilians casualties. While the third condition already requires that harm to civilians be minimized compared to methods of reaching the intended end, it says nothing about the level of risk that the just combatants must take on themselves in pursuing that end. Walzer, unsatisfied with this result, presents the following alternative to the third condition of the Doctrine of Double Effect: “The intention of the actor is good, that is, he aims narrowly at the acceptable effect; the evil effect is not one of his ends, nor is it a means to his ends, and, aware of the evil involved, he seeks to minimize it, accepting costs to himself” (Walzer, 1977: 155). Walzer is advancing a principle with which I agree: just soldiers are sometimes obligated to take risks on themselves in order to minimize harm to civilians. But it is not clear how much this helps the Doctrine of Double Effect; if the requirement to minimize civilian casualties is genuine, then it should not maintain the requirement that all of those casualties be unintended. In the nuclear war case I discuss above, the requirement that civilian casualties be unintended forces us to kill 99,999 more civilians than if we were allowed to kill a single civilian intentionally. The reasoning Walzer employs in support of this second intention seems to better support views without a strict prohibition on intentional civilian casualties.

Walzer presents the doctrine of supreme emergency as an answer to the sort of objection I raise toward the end of the previous section; that is, cases where a principle like the Doctrine of

Double Effect seems to require us to choose the significantly worse of two outcomes. He states the doctrine of supreme emergency as follows:

“A supreme emergency exists when our deepest values and our collective survival are in imminent danger.... There are no moments in human history that are not governed by moral rules; the human world is a world of limitation, and the moral limits are never suspended.... But there are moments when the rules can be and perhaps have to be overridden.” (Walzer, 2004: 34)

Roughly speaking, Walzer is trying to accommodate an intuition that “desperate times call for desperate measures” without abandoning a strong prohibition of those desperate measures. To that end, he criticizes both absolutism and utilitarianism. He criticizes the former for its lack of flexibility in the face of disaster, and the latter for its excessive flexibility. While I will not explore the argument against absolutism (I make similar arguments in the previous section, and am obviously happy to grant the conclusion), I will take a moment to discuss his argument against utilitarianism.

Walzer’s main argument against Utilitarianism seems to be that it is easy to manipulate: “When solidarity collapses, in pure or almost pure adversarial situations – in war, for example – utilitarian calculation is zero-sum, and ‘we’ commonly attach only negative value to ‘their’ utilities” (Walzer, 2004: 39). Since we are making the calculations, and we (supposedly) have an interest in permitting a greater range of actions, we are likely to allow intentional harm to civilians in cases where we should not. But this only seems to tell against our allowing generals to make the calculation; since generals have neither proper authority nor training to decide these questions, it should not be surprising that allowing them to do so would lead to abuse. When a consequentialist of any kind suggests that it is permissible to attack civilians in cases where

doing so leads to greater good for all involved, she is not committed to saying that doing so leads to the greater good when we say it does. Put differently, the consequentialist is as able to criticize the general as the absolutist or rights theorist: she can simply say that the general did not in fact act to promote the greatest good, and so acted immorally.

His argument may succeed, however, if it is understood as aiming at a different claim, that (act) utilitarianism cannot serve as a legal principle governing war. McMahan makes a similar distinction after discussing civilian liability in *Killing in War*. He says:

“If the appropriately limited moral permission were to be legally recognized, the temptation to attack civilians in war is so great that just combatants would inevitably abuse the permission, while unjust combatants, imagining themselves to be just combatants, would liberally avail themselves of it as well. While absolute civilian immunity is false as a moral doctrine, it remains a legal necessity.” (McMahan, 2010: 234-235)

This is probably correct: insofar as the permission to attack civilians only applies to just combatants, it is quite likely that unjust combatants will believe themselves to be just, and so act wrongly. But it is not clear how effective this absolute prohibition on targeting civilians has been; while the question “how many civilians would have been killed if we did not accept this prohibition?” is unanswerable, the massive number of civilian casualties throughout history should, I think, tell against its effectiveness. Further, legal principles that are excessively strict may require us to choose significantly worse outcomes if the better outcome would require violating a law. The question of which state of affairs, either having or not having a legal prohibition on civilian targeting, leads to better consequences is an empirical one that I cannot answer here.

Walzer appears to set an extremely high bar for what counts as a “supreme emergency.” After a discussion about the value of the community, he notes: “dirty hands aren’t permissible (or necessary) when anything less than the ongoingness of the community is at stake, or when the danger that we face is anything less than communal death” (Walzer, 2004: 46). If this is meant as a claim about a moral principle, rather than a legal principle, then I do not see how this can answer the objection I raise toward the end of my discussion of the Doctrine of Double Effect. Perhaps a nuclear attack of the kind I describe would indeed destroy any moral or political community (given that nuclear war is the basis of many post-apocalyptic stories, this does not seem too far-fetched). But we can adjust the case to avoid any sort of community death while retaining the apparent horror of the decision. Suppose that instead of nuclear missiles killing 200,000, that same number would be killed through conventional bombs, and would be killed in relatively smaller groups across a larger number of cities. The intuition that we would be justified in intentionally killing one civilian, as opposed to unintentionally killing 100,000, to prevent this does not seem to change. Much of Walzer’s discussion is aimed at persuading the absolutist that the doctrine of supreme emergency is necessary, and not excessively permissive. Cases like these seem to indicate the opposite: the principle of supreme emergency is not permissive enough, if it requires that 99,999 people must die when killing one would accomplish the same goal.

This issue may simply come down to a difference in intuitions, whether those are intuitions in particular cases or intuitions on the relationship between moral and legal principles. But it seems that a consequentialist is unlikely to be persuaded by Walzer’s argument against utilitarianism; to the extent that someone is not persuaded by that argument, they are also not likely to be persuaded that such a high bar for what constitutes a “supreme emergency” is

necessary. From here, I will discuss two targeting theories that are more flexible than the Doctrine of Double Effect, even together with the doctrine of supreme emergency.

IV. Liability, Intention, and Harm

If what I argue above is correct, we must abandon the Doctrine of Double Effect as a principle for choosing targets in war. Still, we have to account for our intuitions in the above case: when faced with a decision to kill only unjust soldiers or kill only civilians, all else equal, we should choose to kill the soldiers. The principle of distinction that is standard in these discussions holds that any intentional killing of civilians is impermissible; however, if we abandon the role of intention, it is difficult to make sense of this principle. It cannot be impermissible to harm civilians at all in war, as that would make fighting a successful just war impossible. Instead, I think we should criticize the concentration of harm in the above case: rather than claim that civilians cannot be harmed at all, or cannot be attacked intentionally, we can claim that harms should not (in general) be concentrated on civilians when we can concentrate them elsewhere. So, in the above case, because we could choose to concentrate harm on the unjust soldiers, we ought to do so rather than concentrate them on civilians. Why? Because the civilians are less liable to face harms, based both on their responsibility for the current war and their role in the war as it continues. We can look at this liability in two ways. The first, following Jeff McMahan, treats liability to attack as a property that makes defensive attack on the liable permissible. The second treats liability an axiological notion related to desert: to harm someone who is liable on this view would be less bad (or, on some views, good) than to harm someone that is not liable. In this section, I will discuss both sorts of views, beginning with McMahan's account. Both views, as it turns out, will have stricter requirements on the use of drones than the Doctrine of Double Effect.

McMahan holds that the proper criterion for liability to attack is “moral responsibility for an objectively unjustified or wrongful threat” (McMahan, 2010: 38). This requires a bit of

clarification. I will discuss moral responsibility shortly, but as for an objectively unjustified or wrongful threat, McMahan notes:

“In its primary sense, a wrongful threat is a threat of wrongful harm – that is, harm to which the victim is not liable – posed by an action that is objectively wrong.... In a second, broader sense, an individual poses a wrongful threat whenever his action threatens to inflict a wrongful harm¹⁵.” (McMahan, 2010: 38)

Critically, he pairs this sort of justification for harm with a lesser-evil justification (McMahan, 2010: 27)¹⁶, which serves a critical role: it can justify attacking those who are not liable to attack (albeit at much stricter standards) (McMahan, 2010: 199). This would include what McMahan calls Nonresponsible Threats, persons who do pose a threat, but are not morally responsible for it, as well as those who pose no threat at all. For an example of the latter, consider the following case: suppose country A is fighting country B, and a civilian from country C might unknowingly wander into a conflict zone. This civilian clearly does not pose a threat to either side, yet might be harmed unintentionally. The lesser-evil justification would permit some actions that harm this sort of civilian, though it is unlikely to approve of many.

The advantage of this sort of account over views like the Doctrine of Double Effect is clear. McMahan’s account is sensitive to the particular features of whatever case is under consideration. Instead of simply separating individuals into groups of combatants and non-combatants, this liability view can assess civilians and combatants alike on a scale of moral

¹⁵ The “objectively wrong” element of this account refers to actions that either violate the relevant proportionality constraint or are pursuing an unjust war aim. The two senses of “wrongful harm” are relevant to whether just soldiers are liable to attack; McMahan spends most of Chapter 2 of *Killing in War* arguing that they are generally not.

¹⁶ He actually discusses two sorts of lesser-evil justification, but this distinction collapses once the role of intention is removed.

responsibility, and require corresponding actions from just combatants. There are three¹⁷ broad categories of threat on McMahan's view: Culpable Threats, Excused Threats, and Nonresponsible Threats. Nonresponsible Threats are the easiest to handle: they are, on McMahan's description, "a person who without justification threatens to harm someone in a way in which she is not liable, but who is in no way morally responsible for doing so." (McMahan, 2010: 168) One example of this is a symptomless carrier of a disease: this person is harming others who are not liable, but is not morally responsible at all for the fact that she carries a disease yet does not show symptoms (McMahan, 2010: 168). These threats can, of course, be attacked or killed depending on the circumstance, but they are not at all liable, and so require a lesser-evil justification. The benefit from attacking a Nonresponsible Threat must be greater than for other types of threats if the attack is to be justified.

Culpable Threats, on the other hand, are fully liable to attack. McMahan describes them as follows:

"These are people who pose a threat of wrongful harm to others and have neither justification, permission, nor excuse. They may intend the harm they threaten, or the risk they impose, or the threat may arise from action that is reckless or negligent.... Precisely because of the Culpable Threat's full culpability, it can be proportionate to inflict a significantly greater harm on a Culpable Threat if that is necessary to prevent him from inflicting a lesser harm on an innocent victim." (McMahan, 2010: 159)

Culpable Threats are, on this picture, soldiers like the enthusiastic Nazi. He is under no duress, and knows fully what his government is doing, and is eager to fight regardless. It does not seem

¹⁷ Excused Threats are further divided into Partially Excused Threats and Innocent Threats, but I do not think these distinctions are not relevant here.

counterintuitive to say that we ought to stop the enthusiastic Nazi from harming others, even if that requires us to harm the Nazi to a greater degree.

Excused Threats are in between Nonresponsible Threats and Culpable Threats. These threats have some excuse (on McMahan's view, likely some variant of duress or epistemic limitation), and this can either partially or fully excuse their actions. It need not, however, absolve the individual of moral responsibility for the harm done; McMahan offers the case of a soldier under "irresistible duress": this soldier may be acting because his family is being held hostage, or for some similar reason (McMahan, 2010: 162). While the soldier is still responsible for contributing to the unjust war, he is not blameworthy for doing so (let's suppose). The proportionality requirement in this case will be stronger than for Culpable Threats, and just soldiers may be required to take on additional risks to avoid doing excess harm to excused threats, but the exact strength of the requirement will depend on the level of excuse.

This account of threats applies to both unjust combatants and civilians. While it is obviously more difficult for a civilian to be a Culpable Threat than it is for a combatant (if only because it is more difficult for a civilian to be a threat), McMahan offers a pair of cases in which civilians are liable to harm as a result of their status as a threat. The first case is that of Palestinian civilians during the rocket attacks of 2006¹⁸: Palestinians allowed rocket attacks to be launched into Israel from near and among their homes, thereby requiring that any defensive attack by Israel would unintentionally destroy those homes (McMahan, 2010: 220-221). This, McMahan supposes, makes the civilians liable to unintentional harm, though not intentional attack (as the civilians themselves were not attacking Israel). Most of these civilians would likely count as partially excused on McMahan's framework: they may be under some level of duress,

¹⁸ Both McMahan and I are using this case for its structural features, and are not making actual claims about the culpability of Palestinian civilians in 2006.

but by supposition it is not a significant level. McMahan's reliance on the distinction between intended and foreseen harms is something I will argue against shortly.

McMahan's second case, in which the civilians may be liable to intentional attack, proceeds as follows:

“the question of whether the country will resort to war is resolved by plebiscite, the votes in favor of war win overwhelmingly, and an unjust war fought by robots is automatically initiated. Suppose that the victims of the robot attack cannot defeat the robots but can attack the civilian population of the country that has launched the unjust war. Suppose further that it is reasonable for them to believe (and it is also true) that if they kill enough civilians, that will motivate the survivors to rescind the vote and recall the robots.

Intuitively, it seems permissible for the victims of the unjust war to begin to attack and kill the civilians, not because, or only because, this is the lesser evil, but because the civilians have made themselves liable to military attack.” (McMahan, 2010: 227-228)

The culpability of the civilians here has overridden the usually absolute prohibition on intentionally attacking civilians. Critically, however, there are two factors here that impact this liability. First, the robots themselves are clearly neither liable nor capable of being harmed (assuming these robots are not sentient); that is, any military action against them would clearly be a stalling measure until some way of convincing the civilians to end the war could be found. Second, attacking the civilians in this case is stipulated to be effective, where such attacks rarely are, at least when compared to military targets. McMahan acknowledges that the ineffectiveness of attacks on civilians is critically important: he takes that, together with the uneven responsibility of civilians, and that this responsibility is usually quite low (along with a few other factors), to “constitute the real basis of the *moral* immunity of civilians” (McMahan, 2010: 231).

It seems, however, that these considerations can get (a non-absolute) civilian immunity off the ground without appealing to the intentions of the attackers. Next, I will present a case designed to argue against the relevance of intention.

McMahan brings up a particularly interesting implication of his view on liability and excuse. Might there be cases where just soldiers have a duty to minimize the force they use against unjust soldiers, even if that requires taking on additional risks themselves? He believes that there are cases like this, in particular wars involving child soldiers. McMahan says of child soldiers:

“Child soldiers are quite common in some areas of the world.... A child of perhaps 10 years of age is selected and ordered to take a gun and kill his friend, or perhaps one of his parents. The first such child may hesitate or refuse, and is instantly shot dead. Another is then brought forward and given the same order. Those subsequently chosen are less likely to refuse. After a sufficient number of children have been put through this ordeal in full view of the entire village, they are taken away at gunpoint to a camp where they are to be turned into soldiers.” (McMahan, 2010: 199)

Surely if any soldiers are not at all liable to attack, or are at least minimally liable, it is child soldiers. McMahan admits that just soldiers may be required to avoid harming child soldiers unnecessarily: “But when just combatants could use lesser force against child soldiers without seriously compromising their ability to achieve their just aims.... Just combatants should show them mercy... in order to try to allow these already greatly harmed children a chance at life” (McMahan, 2010: 201-202).

We get an interesting result, I think, if we combine features of the child soldier case with that of McMahan’s robot army case. Recall that in that case, there is a civilian population that is

considered culpable for the threat posed by their unjust army. If we agree with McMahan that the civilians in this case are liable to not just unintended harm, but intentional attack, then we can create a bit of a puzzle by replacing the robots with child soldiers. We now have a case where just soldiers are required to restraint themselves in attacking the unjust (child) soldiers, but may intentionally attack the civilians. Surely the permissibility of attacking the civilians did not change when the robots were replaced with child soldiers; if anything, the liability of the civilians has increased, as they are harming these children in addition to pursuing the goals of an unjust war. Suppose we are fighting a just war against this country, and a pilot is presented with two possible bombing raids: (1) the bomber may attack an area in the enemy capital, intentionally killing 1,000 civilians, and ending the war; (2) the bomber may attack a military base, intentionally killing 500 child soldiers and unintentionally killing 500 civilians, and ending the war. Given what we have stipulated about the culpability of the civilians and lack of culpability of the child soldiers, it seems that the pilot ought to pick (1).

If we take cases like this seriously, then we have a peculiar thing to explain: an intentional attack that only kills civilians is easier to justify than one that kills civilians unintentionally while intentionally killing soldiers. Given that intention is supposed to be relevant not to harming those less liable, but instead specifically to harming civilians, this result is strange. If we simply transition to “liability to harm” rather than “liability to (intentionally) attack,” it becomes easier. The latter case is impermissible because we harm those who have a lesser degree of liability to harm, while in the former case only those who are more liable are harmed (they are more responsible, etc.), and we stipulate that the two operations are equally effective. Jettisoning the relevance of intention does not seem to cost much for McMahan’s view: he can still explain what is generally wrong with targeting civilians, and he already

considered the moral responsibility of those harmed as relevant to the permissibility of the action. Rather than maintaining separate proportionality conditions, an account of liability to harm rather than attack can rely on a single proportionality condition, and tune everything else to the responsibility of those harmed in a straightforward manner. The rest of McMahan's liability account seems to be a significant advantage over the Doctrine of Double Effect: it avoids all of the counterexamples I have offered, and does a better job of preserving the idea that most civilians should not be harmed in war. Once the role of intention is removed, we seem to have a perfectly adequate theory of targeting and civilian harm.

Up to this point, I have been discussing liability to harm in terms of a constraint on action, much in the way that the Doctrine of Double Effect is discussed, along with McMahan's "liability to attack" view. This is how the subject of targeting in war has traditionally been understood, in part because of the inherent danger that many see in adopting a consequentialist approach. Peter Olsthoorn, attempting to defend consequentialism, characterizes the worry as the following: "utilitarianism in its paradigmatic form (and consequentialism for more modern varieties)... would supposedly make military expedience override all other concerns" (Olsthoorn, 2011: 81). These concerns are unconvincing if their conclusion is that the proper *moral principle* governing war cannot be consequentialist, as we will see, though they may be more effective if their target is a legal principle governing targeting in war¹⁹. Still, the right consequentialist principle can produce the same results as the deontic principle I have been advancing. This would be a view according to which the desert of the potential subject of harm

¹⁹ For instance, as I mention above, the Doctrine of Double Effect has been defended with claims that even if it is a poor moral principle, it is critical to prevent attacks on civilians, as military commanders are more concerned with military effectiveness and any punishment they might face than with moral concerns. (See McMahan 234-5 for a brief discussion of this point.) Given that some form of the Doctrine of Double Effect has been alive since Aquinas, and that civilian collateral damage (and straightforward war crimes) is disturbingly commonplace, it is not clear that we should accept these claims.

affects the badness of that harm. One such view is Fred Feldman's "Subject's Desert-Adjusted Intrinsic Attitudinal Hedonism" from *Pleasure and the Good Life*:

"i. Every episode of intrinsic attitudinal pleasure is intrinsically good; every episode of intrinsic attitudinal pain is intrinsically bad.

ii. The subject's desert-adjusted intrinsic value of an episode of intrinsic attitudinal pleasure is equal to the amount of pleasure contained in that episode adjusted for subject's desert; the subject's desert-adjusted intrinsic value of an episode of intrinsic attitudinal pain is equal to – (the amount of pain contained in that episode adjusted for the subject's desert).

iii. The intrinsic value of *a possible world* is entirely determined by the subject's desert-adjusted intrinsic values of the episodes of intrinsic attitudinal pleasure and pain contained in *that world*, in such a way that one *world* is intrinsically better than another if and only if the net amount of intrinsic attitudinal pleasure adjusted for the subject's desert in the one is greater than the net amount of that sort of pleasure in the other." (Feldman, 2004: 195)

Essentially, it is worse if good (or innocent) people suffer than if bad people suffer. We can apply this to the case of liability in military actions quite easily: it is worse for a non-combatant to be killed than an unjust combatant (and worse still if the non-combatant is truly innocent, i.e. does not support the unjust combatants). This means it is better if we concentrate harms on unjust combatants than if we fail to do so (or, worse still, concentrate them on civilians). If we have two means of accomplishing a particular military objective (or, instead, two distinct objectives that would make the same amount of progress toward winning the war), and one of them would concentrate harm on civilians, while the other would concentrate harm on unjust

soldiers, we have reason, by an account like this, to prefer the action that concentrates harm on the unjust soldiers. This reason is stronger the more innocent the civilians, and the more culpable the soldier. What does this mean? Consider three cases: Nazi Germany, slave soldiers, and ISIL²⁰.

In the case of Nazi Germany, we have civilian and soldier populations that are both culpable. The soldiers are unjust (perhaps as unjust as have ever existed), and the government was at least initially supported by the civilian population. While harming civilians is still worse than harming soldiers, as the civilians are unlikely to be actively attacking just soldiers, it is less bad to harm civilians in this case, especially if it will lead to less overall harm, less harm to just soldiers, and/or ending the war more quickly. In the case of ISIL, we get the opposite result. The civilians involved are living in villages occupied by ISIL, and are actively being harmed by ISIL's presence (see: beheadings, torture, etc.) The unjust soldiers, on the other hand, are all willing (we'll say) members of a non-governmental group (bracketing considerations of territorial control), and so cannot easily appeal to "just following orders" defenses. In this case, it is *especially bad* to harm civilians in the course of defeating the unjust soldiers. Not only are these civilians maximally innocent, they are currently being harmed by the unjust soldiers around them, and so are, in a sense, especially not deserving of further harm.

The final case, involving slave soldiers, is a bit more complicated. Suppose there is a society²¹ where the general civilian population consists of wealthy landowners operating a democratic government, but the military (except, say, for officers) consists entirely of slaves who

²⁰ I am making some assumptions in both cases that are almost certainly incorrect. For instance, it is unclear how culpable the civilian population was in Germany after Hitler's immediate rise to power; sure, they bear some responsibility for allow/supporting that, but it would take a historian to say how much involvement they had after that point. Similarly, I am sure that some, perhaps many, fight for ISIL against their will. I am simplifying these cases because of how they relate to the principle under discussion, not because these simplifications are historically accurate.

²¹ Those familiar with *A Song of Ice and Fire* or the *Game of Thrones* TV series can imagine Slavers' Bay.

only fight because they are threatened (credibly, let's suppose) with death if they refuse. In this case, we might have a puzzling result: it is *better* to attack the civilian population than the unjust soldiers. Before given a tentative defense of that result, I'd first like to point to one obvious reason why it is not the case. Just as in the case of Nazi Germany, it is the soldiers, not civilians, that are directly opposing the just soldiers, and so harming the soldiers is still the best means of preventing harm to just soldiers and winning the war. But in a case where the progress toward victory is equal, and the potential harm to just soldiers is equal, I do think that civilians ought to receive the greater proportion of harm in a case like this. Not only are the civilians themselves directly responsible for starting and continuing the war (again, assuming fully democratic government among the landowning class we are discussing), they are responsible for the harm facing the unjust soldiers as well; we have stipulated that the soldiers only fight because they are compelled to do so. Here, we face a situation similar to that of ISIL: we have reason to avoid harming the innocent group that has already been harmed. Thankfully, cases like this are rare (and perhaps extinct), and it seems likely that such a war would be short: upon discovering the just soldiers can simply circumvent the slave army and attack the slave-owners, the war would likely end rather quickly.

It is important to mention that there might be several versions of a principle like this. The one I have discussed involves affecting the level of badness involved in harming different groups. Others, like Moore's, might consider the harming of some groups positively good, while the harming of others is especially bad. This view has some plausibility: there may be people that are bad enough that the world is better off for harming them. As an example²², consider a world in which Hitler survived World War II. A branch of this world in which Hitler lives out the rest of his days in comfort does seem somehow worse than one in which he is punished. To modify

²² Thanks to Alastair Norcross for suggesting this.

Ross's Two Worlds argument, a world in which there is a mix of good and bad people, and the good are happy while the bad are punished does not seem obviously worse than one in which there are good and bad people and no one is punished. One might argue that there is a lack of fitting response to the badness of the people in the world in which no one is punished.

If we translate this Moorean view to war, however, we may encounter more difficult examples. If we consider the case of slave soldiers, or the child soldier case above, we may find we have civilians that we *ought to* attack, not simply instead of attacking soldiers, but that we ought to attack regardless. Still, it is unlikely, for the reasons McMahan gives, and I discuss above, that any civilians should be liable or deserving to that extreme degree. And if such horrible civilians were to exist, it may not be so counterintuitive that harming them is good. Though I will remain with Feldman's view here, whatever difference that makes, these other views seem similarly plausible.

I will mention a couple of concerns before moving forward. First, I do not have any particularly robust conception of desert in mind here. Roughly, I imagine desert in these cases will roughly track the "liability" I have discussed above: if someone is not liable to harm, then she does not deserve that harm, and vice versa²³. Second, it might be that a full account of desert-adjusted consequentialism would have results that contradict my claims about the particular cases here, or that there are sound arguments against these sort of views. I am not concerned with that here. Instead, I am concerned with presenting what I take to be an interesting alternative to the deontic principle I advance above.

There is one more question about this view that needs addressing before we can return to the topic of drones: how is the risk to just combatants to be weighed against harm to civilians or

²³ McMahan mentions in a few places that he means to distinguish liability from desert; the idea of desert he has in mind is much more narrow than what I am using here.

unjust combatants. That is, if harming innocent civilians is worse than harming unjust combatants, where does harm to just combatants fit? The answer, I think, depends on the liability of the civilians. As McMahan notes, the just combatant is not merely innocent, but is (by supposition) pursuing the greater good:

“Because the action by the just combatants is, by hypothesis, justified, it must be proportionate in the wide sense, which means its good effects outweigh the harm to the civilians. Killing the just combatants, thereby thwarting their mission, therefore cannot be the lesser evil.... For the relevant bad effects of attacking them now include not only the failure of their just mission, which alone outweighs the deaths of the innocent civilians, but also all the harms they suffer.” (McMahan, 2010: 45)

If the harm to just combatants threatens to cause their mission to fail, then it is clearly much worse than it would otherwise be. In cases where it would not lead to mission failure, however, we get something similar to the case of truly innocent civilians, like in the ISIL case above. In each case, there is a factor that amplifies the badness of the harm done: in the case of just soldiers, it is their pursuit of a just cause, where in the case of the civilians in ISIL, it is the fact that they are already being harmed by unjust soldiers. In each case, this factor seems to make the relevant group less deserving of harm.

In cases where the harm done to the just soldiers results in the failure of their mission, it seems, then, that this harm would, effectively, “trump” harm done to either civilians (of any type) or unjust soldiers. In cases where that harm would not cause the failure of the mission, or where it is merely the risk of harm, we have a more interesting result. Cases like ISIL, I think, can generate obligations for just soldiers to increase risk of harm to themselves to prevent causing harm to innocent civilians. The badness that results from harming innocent civilians is

greater than that resulting from the increased risk to just soldiers, so long as the just cause is not itself threatened. By contrast, cases like Nazi Germany do not have this result – the culpability of the civilians in that case lessens the badness of harming them such that the just soldiers are likely not required to take on additional risks. Cases like Slave Soldiers might even generate an obligation in the opposite direction: just soldiers might be obligated to avoid harming unjust soldiers when possible, even if that comes at the cost of additional risk to the just soldiers²⁴. This matches the results of McMahan’s liability view, as I discuss above: in many cases involving civilians, due to their low level of liability, and in cases where the unjust soldiers are sufficiently excused, just soldiers may be required to take greater risk on themselves in order to minimize harm to the respective group.

Now, finally, we can return to the case of drones. If what I argue earlier is true, and drones cause a significant amount of psychological harm on nearby civilians, then just soldiers might have an obligation to avoid causing that harm, even if alternative strategies place the risk of harm on just soldiers. Using drones to fight terrorism, for instance, is likely impermissible in many cases, if those cases are structurally similar to ISIL (and I suspect they are). In cases where the terrorists have broad local support, this obligation would be less likely. These cases would be more akin to Nazi Germany, where the culpability of the civilians renders them more liable to harm than the just soldiers. This is a good result: claims that, for instance, we are not permitted to use drones in traditional war are strange, precisely because it seems that Strawser’s argument fits those cases best. We can pursue our military aims with equal (or greater) effectiveness when compared to alternative strategies, with close to no cost to our soldiers. The concerns about

²⁴ This is, of course, entirely dependent on the circumstances, including the effectiveness of the unjust soldiers. I am not suggesting that just soldiers must sacrifice themselves for the sake of unjust soldiers, but that, given the minimal culpability of the unjust soldiers, that just soldiers should not exterminate them for a minimal amount of additional security.

psychological harm seem most relevant to the permissibility of drone use exactly in those cases where the civilians are completely undeserving of that harm²⁵. My modified version of McMahan's principle of liability yields the same result. In cases where the civilians are not at all liable to harm, we cannot harm them if there is an alternative strategy than can sufficiently accomplish our mission, even if that results in greater risk to our soldiers. Since these alternative strategies seem to exist (special forces, and so on), we are obligated to employ them rather than use drones. In cases of conventional war, akin to Nazi Germany, the increased liability of civilians may remove this requirement. Still, even in these cases, the allowed harm would be bound by the level of civilian liability, and there may well be cases where the use of drone is impermissible (for instance, hovering drones over a large enemy city with few or no military targets). We have, it seems, two plausible accounts of distributing harm in war, and both agree that the psychological harm caused by drones renders their use impermissible.

²⁵ This is extremely relevant to our drone policy, as this comprises nearly all of the cases in which drones are used. Since the United States' drone program has been expanded under President Obama, the US has largely withdrawn from conventional war.

V. Conclusion: the Proper Role of Drones

If everything I have said so far is true, we must place significant restrictions on the use of drones. Using drones in counterterrorist operations is almost certainly impermissible, at least as we currently conduct those operations. These operations too frequently occur in and around civilian areas, leading to the psychological harm discussed toward the beginning of this paper. Further compounding this problem is the international perception of American drone use. A Pew Research Center survey from spring 2014 lists only four countries as having a greater level of approval than disapproval of American drone use: the United States, Nigeria, Kenya, and Israel (“Widespread Opposition to Drones,” 2014). Much of this international opposition can, I think, be attributed to the psychological harm caused by drones. As noted above, civilian deaths are significantly lower when drones are used than in conventional warfare; while these casualties draw more attention now than they might have in the past, I do not think that this alone can account for international opposition to drone use. In any case, the international perception of drones seems to increase the level of psychological harm they may cause on civilians: someone may see a drone in the sky and be afraid because constantly hear criticism of American drone use. Restricting our use of drones, however, may lead to the opposite effect: if the international perception of our actions reverses²⁶, civilians are likely to have a positive reaction to seeing drones, rather than the reaction they currently have (detailed toward the beginning of this paper). As a result, we have somewhat of a feedback loop: as we use drones impermissibly, people begin to fear them in more cases, further restricting how we are permitted to use them, and so on. Severe restrictions are required, as I have argued, but there is still a significant role that drones can serve in military policy. I will describe some of these uses here.

²⁶ Of course, some of the international opposition to US drone policy is politically motivated, and so is unlikely to respond to any actual change in our policy. But it seems overly cynical to attribute too much of the opposition to this.

First, drones can have a significantly expanded role in conventional warfare. Let us return to the case of Nazi Germany. I argue above that drone use is permitted in this case because the civilians are more liable, or more deserving, of any harm they might receive. There is a further reason, however, that supports the use of drones in cases like this, or in other cases of conventional warfare (involving, say, a less straightforwardly evil enemy). Civilian areas near the front lines of any large-scale war will already be suffering from the sort of psychological harm I attribute to drone use: there is a constant risk of physical harm, uncertainty as to whether the area will be attacked, constant visual and auditory reminders of combat, and so on²⁷. In such a case, it is clear that drones are unlikely to cause any further psychological harm than the civilians are already suffering; the drones would act merely as more precise artillery. In this case, the greater precision of drones with respect to physical harm gives us reason to prefer them to other weapons or strategies. Further, the increased liability of the civilians in this case lessens any requirement for just soldiers to take further risk on themselves; drones, the use of which effectively removes any risk from soldiers on the side using them, are no longer getting in the way of that requirement. The use of drones in conventional war would fit Strawser's argument completely, and so there is a strong case to be made that we are obligated to use them in many cases.

There may even be cases of counterterrorism in which the use of drones is permissible. While I have argued extensively against the permissibility of drone use in the "standard" counterterrorism case, there are other cases that do not involve civilians suffering psychological harm. If we discover a terrorist training camp that is away from civilian areas, for instance, there would be nothing wrong with using drones to attack that area: there are, by stipulation, no civilians to be harmed, and the terrorists themselves are liable to harm. This sort of case might be

²⁷ Imagine a village in an area under artillery fire.

increasingly uncommon, particularly if the United States were to adopt the sort of restrictions on drone use that I suggest here, but any cases like this would be unproblematic. A perhaps more helpful use for drones would be to survey, and attack if possible, mountain caves and other similarly remote locations. Here, if there is no civilian presence, drones can monitor the area for terrorist activity and strike when warranted.

Finally, there are many non-military applications for drones that would not be affected at all by the arguments here. I will not explore these in detail, but the use of drones for humanitarian purposes would obviously not be made impermissible by arguments against the use of drones in combat. What I hope to have shown, at least, is that while drones open up exciting possibilities in warfare, they also bring with them new harms. If we reject the role of intention in assessing permissibility, as I argue we should, then it becomes much more difficult to justify inflicting these harms on civilians who are not liable or deserving of harm. Therefore, we ought to restrict the use of drones in these cases, despite the military advantages they bring.

Works Cited

- Feldman, Fred. *Pleasure and the Good Life: Concerning the Nature, Varieties, and Plausibility of Hedonism*. Oxford: Oxford University Press, 2004. Print.
- “Living Under Drones: Death, Injury and Trauma to Civilians from US Drone Practices in Pakistan.” *Living Under Drones*. Stanford Law School, 2012. Web. 15 May 2015.
- McMahan, Jeff. *Killing in War*. Oxford: Oxford University Press, 2009.
- McMahan, Jeff. “Intention, Permissibility, Terrorism, and War.” *Philosophical Perspectives* 23: 2009. Web. 15 May 2015.
- Norcross, Alastair J. “The Road to Hell.” Draft. 2011. (Cited with permission.)
- Olsthoorn, Peter. “Intentions and Consequences in Military Ethics.” *Journal of Military Ethics* 10.2: 2011. Web. 15 May 2015.
- Plaw, Avery. “Counting the Dead: The Proportionality of Predation in Pakistan.” *Killing by Remote Control: Ethics of an Unmanned Military*. Ed. Bradley Jay Strawser and Jeff McMahan. Oxford: Oxford University Press, 2013. 126-153.
- Rhode, David. “Reuters magazine: the drone wars.” *Reuters.com* Thomson Reuters, 26 January 2012. Web. 15 May 2015.
- Strawser, Bradley Jay. “Moral Predators: The Duty to Employ Uninhabited Aerial Vehicles.” *Journal of Military Ethics* 9.4 (2010): 342-368.
- Walzer, Michael. “Emergency Ethics.” *Arguing About War*. New Haven: Yale University Press, 2004. Print. 33-50.
- Walzer, Michael. *Just and Unjust Wars*. New York: Basic Books, 1977. Print.
- “Widespread Opposition to Drones.” *Spring 2014 Global Attitudes Survey*. Pew Research Center, 2014. Web. 15 May 2015.