

- Hurka, T. (2008), "Proportionality and Necessity," in May, L. (ed.), *War: Essays in Political Philosophy* (New York: Cambridge University Press).
- Khan, P.W. (2002), "The Paradox of Riskless Warfare," *Philosophy and Public Policy Quarterly*, 22:3, 2–8.
- Killmister, S. (2008), "Remote Weaponry: The Ethical Implications," *Journal of Applied Philosophy*.
- Kleinig, J. (2010), *Ethics and Criminal Justice: An Introduction* (New York: Cambridge University Press).
- Lauder, S. (2010), "Trial of Bionic Eyes within Three Years," ABC News (March 30. Available at <http://www.abc.net.au/news/stories/2010/03/30/2860256.htm>, accessed August 20, 2010).
- Licklider, J. (1960), "Man-Computer Symbiosis," HFE-1 *IEEE Transactions on Human Factors in Electronics*, 4, 4–11.
- McMahan, J. (2009), *Killing In War* (Oxford: Oxford University Press).
- Rodin, D. (2002), *War and Self-Defense* (Oxford: Oxford University Press).
- Scerri, P., Von Gonten, T., Fudge, G., Owens, S., and Sycara, K. (2008), "Transitioning multiagent technology to UAV applications," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track* (Estoril, Portugal), 89–96.
- Scheutz, M., Schermerhorn, P. and Bauer, P. (2005), "The utility of heterogeneous swarms of simple UAVs with limited sensory capacity in detection and tracking tasks," *Swarm Intelligence Symposium*, SIS 2005. Proceedings 2005 IEEE (June 2005), 257–64.
- Santhanam, G., Ryu, S.I., Yu, B.M., Afshar, A. and Shenoy, K.V. (2006), "A high-performance brain-computer interface," *Nature*, 422, 195–8.
- Sparrow, R. (2009), "Predators or plowshares? Arms control of robotic weapons," *Technology and Society Magazine*, IEEE, 28:1 (Spring), 25–9.
- Sparrow, R. (2007), "Killer Robots," *Journal of Applied Philosophy*, 24:1, 62–77.
- Tybaji, A. (2007), "Unique problems associated with UAV employment," *Flying Safety* (May 2007).
- Walzer, M. (2000), *Just and Unjust Wars: A Moral Argument With Historical Illustrations* (New York: Basic Books).

Chapter 7

Robotic Weapons and the Future of War

Robert Sparrow

Introduction

Robotic weapons are likely to play a central role in future wars (Featherstone 2007; Graham 2006; Hanley 2007; Hockmuth 2007; Peterson 2005; Scarborough 2005; Singer 2009).¹ The perceived success of the Predator uninhabited combat aerial vehicle (UCAV) in Iraq and Afghanistan has resulted in military forces around the world rushing to develop and field unmanned systems (UMS) (Hockmuth 2007; Office of the Secretary of Defense 2005). Existing UMS are, for the most part, operated by a remote operator but some also make use of sophisticated artificial intelligence technology in order to provide the weapon with some capacity for autonomous action (Excell 2007; Braybrook 2007). In the future, the scope of autonomous action allowed to UMS may increase, with the possibility that Autonomous Weapon Systems (AWS), capable of being deployed to achieve particular military objectives without human supervision, will eventually be developed (Office of the Secretary of Defense 2005, p. 52; Singer 2009, pp. 123–34).

It is clear that the use of UMS in warfare will generate a number of challenging ethical issues—more than can be dealt with in a single chapter.² I have written about the difficulties involved in attributing responsibility for deaths resulting from the actions of autonomous weapon systems (Sparrow 2007a), and the possibility that robotic weapons will lower the threshold of conflict, elsewhere (Sparrow 2009a). I have also addressed the issues that may confront engineers and other researchers involved in the design of these systems (Sparrow 2009b). In this chapter I want to examine a number of issues raised by the development and deployment of robotic weapons that should be of special interest to those writing and thinking about the ethics of war more generally. That is, I have here chosen to discuss those issues

¹ The research for this chapter was supported under the Australian Research Council's Discovery Projects funding scheme (project DP0770545). The views expressed herein are those of the authors and are not necessarily those of the Australian Research Council. I would like to thank Neil McKinnon for assistance with research for this chapter. I would also like to thank Jessica Wolfendale, Tony Coady, Linda Barclay, Jim Sparrow, and Ron Arkin for comments and discussion that have improved this chapter.

² Other discussions include: Asaro 2008; Krishnan 2009; Borenstein 2008; Sharkey 2009; Singer 2009.

that raise the most pressing philosophical problems and/or where the impacts of robotic weapons on the nature of armed conflict are likely most profound.

With any discussion of the ethical issues raised by a given technology, there inevitably arises the question of whether the issues are in fact unique to the technology under discussion or whether they are associated with the technology's predecessors. Is the "new" technology really that novel or is it just an extension of existing and historical technologies? (Sparrow 2007b). This question arises with especial force regarding military robotics because there is a long history of the use of machines in warfare. Both of the features of robots that are responsible for the majority of qualms about their military applications—that they allow "killing at a distance" and that the operations of the machine may determine who lives or dies—are shared with other technological systems. It might therefore be argued that military robotics is undeserving of being singled out for focused ethical examination.

For current purposes, I wish to remain agnostic on the question of the extent to which the ethical issues raised by robotics are unique to robots. While many of the issues I discuss below are familiar, others are less so. In particular, some of the issues raised by the (hypothetical) development of *autonomous* weapon systems seem to have relatively few precedents in existing weapons. More generally, recent and likely future developments in robotics exacerbate and bring into stark relief issues that arise elsewhere in less dramatic forms. In any case, ethical issues do not need to be novel to be important. Given that robotic weapons are coming into widespread use for the first time now, it is appropriate—indeed, necessary—to give whatever ethical issues they raise serious consideration. It is to this task I now turn.

The Ethics of Killing-at-a-distance

Perhaps the first issue to arise in any discussion of robotic weapons is the ethics of killing at a distance. The pilots who operate Predator and Reaper are safely ensconced thousands of kilometres away from the people they kill. The drones they operate fly so high and so quietly that the people that they target often have no idea that they are there until their missiles arrive. It is hard *not* to wonder about the ethics of killing in this fashion. In particular, is hard not to think that by distancing warfighters from the consequences of their actions, both geographically and emotionally, UMS will make it more likely that the operators of robotic weapons will make the decision to kill too lightly (Bender 2005; Ulin 2005).

Ultimately, whether the use of UMS will contribute to more or less ethical behaviour by their operators is an empirical matter that can only be resolved by careful study over an extended period of time. However, there are number of reasons to think that robotic weapons are less problematic in this regard than first appears.

To begin with, any analysis of the impact of robotic weapons on the behaviour of those operating them must keep in mind the character of the weapons these systems replace. Warfare has involved "killing at a distance" ever since the invention of the sling and the spear. Modern warfare includes many weapon systems that allow their operators to kill people they have never met and that distance those that operate them from the destruction that they wreak (Sullivan 2006). The operators of UAVs, at least, are *better* situated to learn about the consequences of their actions than, for instance, artillery gunners, bombardiers, or other warfighters who kill at a distance using existing weapons systems. The pilots who fly UAVs are able—indeed, are usually required—to conduct a battlefield damage assessment which exposes them to (some) knowledge of the consequences of their actions.³ One might therefore hope that these pilots will take the decision to kill at least as seriously as their comrades. Other remotely operated systems (ground vehicles, submersibles, surface vehicles, etc) may not offer as good a perspective on the battlespace as UAVs but will still usually do better than comparable systems without operators. It is also worth noting that the telemetry from remotely operated weapon systems will usually be recorded and be available for review, which may also serve as a significant incentive to obey the laws of war.

It must also be acknowledged that close proximity between enemies is far from being a panacea against illegitimate killing. Some of the worst atrocities of modern times have been carried out at close quarters by men armed with rifles and machetes. Any attempts to evaluate the impact of UMS on the extent to which those fighting wars are willing to kill will need to be made in the context of a realistic appraisal of ways in which people behave in more familiar and less high-tech forms of conflict (Arkin 2009, 30–6). If the appropriate comparison is Kosovo or Rwanda then the argument that UMS will make killing more likely may be hard to establish.⁴

³ However, whether the sensor systems of UAVs are capable of communicating the reality of these consequences in such a way that knowledge of this reality might inform the future moral character of those who kill using these systems, is less clear. While video taken from 15,000 feet may be adequate to determine whether tanks or artillery have been destroyed and people killed, it may not be enough to convey the *moral* reality of killing (Cummings 2004). The sorts of experiences that are essential to understanding what it means to kill may be too "embodied" and subtle to be transmitted via video (Mander 1978). For further discussion, see Sparrow (2009b).

⁴ My thanks to Linda Barclay for encouraging me to admit the force of this objection.

How will Robots Transform the Requirements of Discrimination and Proportionality?

An issue which has received some useful discussion in the literature already is the way in which the development of UMS may affect the proper interpretation of the principle of discrimination within the law of war (Meilinger 2001). This principle, which is a central part of the just war doctrine of *jus in bello*, requires that warfighters not directly target non-combatants and also take reasonable steps to avoid non-combatant casualties (Lee 2004). What is reasonable in this context depends in part on the technical means available to the warfighter (Dunlap 1999, Schmitt 2005). UMS can greatly increase the precision with which weapons can be targeted by making available real-time targeting information (Kenyon 2006, Meilinger 2001; Tuttle 2003). They may also be capable of delivering weapons in such a fashion as to ensure maximum accuracy in circumstances where it would be too risky for a manned system to attempt to do so. The more it becomes possible to identify and destroy specific targets using UMS, the more commanding officers may be obligated to make use of these systems. An unintended consequence of the development of these more versatile systems may be that commanders eventually face new restrictions in their use of older generations of weaponry.

The requirements of proportionality may also be affected by the development of UMS. Under *jus in bello* warfighters are only to use force proportionate to legitimate military goals. Again, what is proportionate use of force will depend, in part, on the range of possible applications of force that are available to achieve the required goal. In so far as UMS allow warfighters to apply destructive force more precisely they will also allow them to use a smaller amount of force. The destruction involved in the use of other weapon systems may eventually come to appear disproportionate by comparison.

The development of UMS may affect the requirements of proportionality via another route. The degree of force that it is permissible to use depends on the nature and importance of the military goal that is being pursued. When warfighters' lives are at stake, it is acceptable to use a large amount of destructive power to preserve them. However, in operations involving (only) UMS, enemy combatants will pose no threat to the lives of their operators. Even if this does not entirely undercut the justification for the use of lethal force, it may nevertheless affect our judgements about the level of force that is proportionate in response to an attack on a UMS. Obviously, the degree to which the absence of a threat to the operator diminishes the amount of violence that constitutes a proportionate response will depend on the particular circumstances in which an operator of the UMS might be considering using lethal force and the nature of the conflict in which the system is engaged. While enemy combatants may not pose an urgent threat to the operators of the UMS, they may pose such a threat to other warfighters. Similarly, while the destruction of a UMS might not harm its operators, it may threaten lives elsewhere in the battlespace or at another time when the UMS would otherwise have played a role in reducing the threat to

combatants. If we think of enemy warfighters in general as posing a threat—as I will discuss further below—then the operators of UMS may be justified in using whatever amount of force is necessary to kill them where-ever they encounter them. However, if our sense of a proportionate response is linked to the threat posed by a particular unit in particular circumstances, then the proportionate response from a UMS to a military threat may be significantly different to that of a manned system.

Importantly, the development of UMS also opens up the possibility that *new* targets may become available for legitimate attack (Hambling 2007). In so far as they make possible more discriminating and proportionate uses of lethal force, UMS may allow warfighters to attack targets that previously could not be attacked without risking excessive casualties (Arkin 2009; Dunlap 1999; Meilinger 2001; Metz 2000; Office of the Secretary of Defense 2005; Schmitt 2005). Of course, this also means that military forces may attack *more* targets, which in turn may have implications for the likelihood that military campaigns will produce civilian casualties.

The implications of the development of UMS for the requirements of the principles of discrimination and proportionality are thus more complex than first appears. It is clear, however, that the proper interpretation of the principles of discrimination and proportionality will be affected by the evolution of robotic weapons. Changes to the interpretation of these principles may have significant implications for the future of war.

Autonomous Weapon Systems and Jus in Bello

There is another important question that arises about the principles of discrimination and proportionality if *autonomous* weapon systems come into widespread use. That is, if decisions about targeting and the use of lethal force are eventually handed over to machines.

Initial discussions of robotic weapons by those advocating their adoption were careful to insist that the decision to take human life would never be made by a machine. While a robot might identify and track a target, the decision to attack it would always be made by a human being. Thus, there would always be a human being “in the loop.” More recent discussions have come to endorse a significantly looser idea of human control over the operations of the weapon system as a whole rather than each targeting decision it makes—a human “on the loop” rather than “in the loop” (Arkin 2009; Singer 2009, pp. 123–34). With the benefit of hindsight, it is reasonably clear that the previous insistence that machines will never be allowed to make a decision about targeting human beings was, at the very least, naive and probably disingenuous (Singer 2009, pp. 123–34). The logic of the dynamic that has produced these weapons pushes towards the development of autonomous weapon systems (Adams 2001). Now that it has become possible to take the pilot out of the weapon platform, the question naturally arises whether it might be possible to take them out of the system altogether. The tempo of battle in air combat is already

so high that it is difficult for human pilots to make decisions in the time available to them. Consequently, humans already rely upon automated systems such as threat assessment displays to make life or death decisions. Eventually the tempo of battle may become so high that it is impossible for human beings to compete with an automated system that has the authority to fire weapons; at this point, the temptation to hand over control of the entire platform to an automated system will be almost irresistible. Moreover, the communication infrastructure necessary to maintain the link between the operator and the weapon is an obvious weak point in the operations of the existing drones. Competition for satellite bandwidth already places significant limits on the number of drones that can operate over Iraq and Afghanistan. An enemy with the capacity to do so would almost certainly try to destroy or disrupt this communication infrastructure in the opening stages of a conflict. Again, the obvious response is to eliminate the need for a human “in the loop” altogether. Given these pressures, it is important that we think seriously about whether it is possible for robots to obey the dictates of *jus in bello*.

Ron Arkin, Professor of Robotics at Georgia Tech, is already advocating the development of weapon systems capable of operating in a fully autonomous mode (Arkin 2009). Designing robots that are capable of discriminating between legitimate and illegitimate targets under the laws of war involves formidable technical challenges. It is difficult enough for a robot to distinguish a person from their surroundings, let alone a person in uniform from a person not in uniform or—most importantly—a combatant from a noncombatant. Nevertheless, Arkin argues that engineers should take up this challenge, even to the point of aiming to develop an “ethical governor” that can sit within the decision-making processes of military robots in order to allow these machines to distinguish between ethical and unethical killing (Arkin 2009, pp.127–53).

I am inclined to think that this project underestimates the extent to which the proper interpretation of the principles of *jus in bello* is context dependent. Determining the application of concepts that are central to the principles of *jus in bello*, such as “proportionality,” “threat,” and “reasonableness,” requires judgement at an extremely high level of abstraction. Whether or not an armed tribesman in Afghanistan, for instance, is a legitimate target for attack by US forces in the country may depend upon his location, his past movements, his tribal and political affiliations (if known), the history of the local area, and the disposition of other forces and civilians nearby. For a machine to gather the relevant information, let alone represent or process it, would require a computer of unprecedented power and sophistication. Even in more conventional warfare, not all enemy soldiers in uniform are combatants—as they may have indicated their desire to surrender or be so incapacitated by wounds as to no longer represent a threat. Indeed, morally speaking, even uniformed soldiers holding weapons may not be legitimate targets in some circumstances, for instance, if their deaths would be disproportionate to the military goal that would be achieved by killing them. The extent to which judgements about discrimination and proportionality are sensitive to context makes it hard to imagine how any system short of a genuine

artificial intelligence might even begin to apply the principles of *jus in bello*.⁵ For this reason, I remain cynical about the likelihood that automated systems will ever be good enough to be entrusted with decisions about who to attack.

Arkin has a powerful reply available to him at this juncture, which is to point out that it is not necessary to achieve 100 percent accuracy in discriminating between legitimate and illegitimate targets but only to achieve a level of discrimination comparable to that of human warfighters. Moreover, as Arkin (2009, pp. 30–6) observes, human beings are not that good at applying the principles of *jus in bello*. In practice, human beings often fail to discriminate between legitimate and illegitimate targets in wartime, for at least two reasons. First, discrimination is difficult. Decisions often need to be made quickly, under considerable pressure, in chaotic circumstances, and often on the basis of limited information and resources. Little wonder, then, that even the most conscientious of warfighters sometimes make mistakes. Second, even where they are capable of distinguishing between these categories, they may be insufficiently committed to making the distinction. Obeying the dictates of the principles of *jus in bello* will sometimes require significant personal sacrifice from warfighters in terms of the risks to which they must subject themselves. If they hesitate too long before firing on a target or fail to fire on a target about which they are unsure, this may lead to their being killed. Similarly, if they refrain from firing on enemy combatants for fear of causing disproportionate civilian casualties, this may allow the enemy to escape or to continue firing on them. There are also significant social and psychological pressures on warfighters, which may encourage them to come to view all of the members of an enemy political, ethnic, or racial group as a threat to them or otherwise undeserving of concern and respect. All too often, these factors combine to produce an attitude that might be expressed colloquially as “kill ’em all, let God sort them out.”

Given that human beings tend only to apply the rules of *jus in bello* imperfectly in practice, it is not inconceivable that in some domains automated systems will prove more reliable than human warfighters at distinguishing legitimate from illegitimate targets. This is even likely in applications, such as air combat or missile defence, wherein the tempo of battle is such that human beings *cannot* make decisions about firing (or not) in time to be effective (Adams 2001).

However, this way of understanding the principle of discrimination treats it as though it were merely a technical challenge—a problem of pattern

⁵ In fairness to Arkin, he is aware of this problem and suggests that autonomous weapons should be designed to call for assistance from a human supervisor when they encounter a situation that requires judgement of this sort (Arkin 2009, 177–209). This presumes, of course, that reliably determining when one is confronting an ethical dilemma is not also a task requiring a high level of context sensitive abstract reasoning. Moreover, any need for access to human supervision substantially undermines the purported advantages of autonomous weapon systems and may be expected to generate real tension with the imperatives that are driving the development of such systems.

recognition—of the sort that roboticists routinely confront. There may be more to the principle of discrimination than this.

In a famous and influential article on the moral foundations of the principles of *jus in bello*, Thomas Nagel (1972) argued that the force of these principles could only be explained by the idea that they are founded in absolutist moral reasoning. Consequentialist moral philosophy is shaky grounds upon which to found the injunctions of just war theory, as it is all too easy to imagine circumstances in which a clearheaded calculation of consequences speaks against respecting the rules of war. Thus Nagel offers an account of the key injunctions of *jus in bello* by way of a (essentially Kantian) principle of respect for the moral humanity of those involved in war. He argues that even during wartime it is essential that we acknowledge the personhood of those with whom we interact and that,

whatever one does to another person intentionally must be aimed at him as a subject, with the intention that he receive it as a subject. It should manifest an attitude to him rather than just to the situation, and he should be able to recognise it and identify himself as its object. (p. 136)

Another way of putting this is that we must maintain an “interpersonal” relationship with other human beings even during wartime. Obviously, if this principle is to serve as a guide to the ethics of war—rather than a prohibition against it—the decision to take another person’s life must be compatible with such a relationship (Nagel 1972, p. 138).

Nagel’s discussion suggests another interpretation of the principle of discrimination, in terms of what is required to extend people the sort of respect that is appropriate to making the decision to end someone’s life. Discrimination is the act of establishing the interpersonal relationship with a possible target of attack and responding to relevant features of this person, which may or may not justify attacking them with lethal force. If this account of the nature of discrimination is correct then merely showing that a weapon system reliably attacks only legitimate targets will *not* be enough to argue that it is discriminating in the appropriate sense.

Indeed—at first sight at least—it appears that autonomous weapons would *not* be capable of discrimination on this account because they are not capable of establishing the appropriate interpersonal relationship with those they target. When a robot does the killing there is no one who cares whether those killed are combatants or non-combatants or what it is that combatants are doing that makes them legitimate targets. According to this way of thinking, placing the lives of our enemies at the mercy of a machine seems to neglect and devalue their humanity in quite a profound way; we do not even grant them the courtesy of deciding whether they are worth killing.

However, this presumes that it is the robot that should be manifesting the respect. Unless the machine is morally autonomous—which would require it, at

the very least (Sparrow 2004), to be a genuine artificial intelligence—it might be argued that the proper place to look for the required attitude is rather in the person who ordered the deployment of the AWS. Many existing weapons systems—including those that AWS might plausibly replace—seem to involve only a very abstract and mediated relationship between those who use them and those they are trying to kill. If there can be an “interpersonal” relationship between a bombardier and the enemy combatants they target thousands of feet below, it is not clear why a similar relationship could not exist between the person who orders the deployment of an autonomous weapon system and the people who are killed by that system. We might, then, think of an autonomous weapon as being merely a better (smarter?) “smart bomb.” Given the choice, enemy civilians might well prefer that they face robotic weapons rather than bombing from a B52 or even than gunfire from nervous 19-year-old Marines. Insofar as the principles of *jus in bello* sometimes also protect the lives of enemy combatants it is even possible that enemy combatants would prefer to be fighting robotic weapons that reliably implement these principles rather than human soldiers or other technologies that do not. Thus, it seems arguable that the use of autonomous weapons *is* compatible with a respect for the humanity of those being killed.

How we think about the morality of autonomous weapons will depend in part, then, on how plausible it is to understand them as simply a means whereby the person who commands their use kills people. If we think of AWS as transmitting their commanding officer’s intentions, then there will be no question as to whether the robot itself is discriminating in the right way to satisfy the principle of discrimination. However, to the extent that we think of the machine itself as choosing its targets, then even a AWS that is better than a human being at selecting the “right” targets may not be capable of discrimination according to the analysis I have developed here. Which of these interpretations is more plausible is, I think, likely to become one of the key questions about the ethics of the deployment of autonomous weapon systems.

There is one further factor to consider. As Nagel observed in his original paper, how we express respect is, in part, a matter of convention. Thus, for instance, what counts as desecrating a corpse will be determined by conventional social meanings and practices, etc. Thus, prevailing social understandings about what is required to respect one’s enemy have some status simply by virtue of being widely shared. Moreover, the military’s reluctance to admit that machines will eventually be given the capacity to choose their own targets and decide when to release their weapons reflects, I think, their awareness of the strength of the public’s feeling about the idea that machine should have the right to decide who lives or dies. Giving an autonomous weapon systems the power to kill seems a bit too much like setting a mousetrap for human beings; to do so would be to treat our enemies like vermin.

The requirements of convention, then, may well swing the balance of considerations towards the conclusion that even (hypothetical) autonomous weapon systems that are capable of reliably distinguishing between legitimate and illegitimate targets under the laws of war would fail to extend the appropriate respect

to those they might kill. Of course, such a “conventional” account is vulnerable to the possibility that social understandings may change, especially as the public gets more used to the idea of robotic and autonomous weapons. Moreover, it must be acknowledged that conventions are shallow foundations for moral principles that must stand up to the temptation to bend or break them that almost inevitably occurs in war. Thus while I think the balance of considerations favors the conclusion that the use of autonomous weapon systems would be unethical by virtue of its failure to respect the personhood of their targets, this conclusion must remain a tentative and provisional one, pending further philosophical discussion and investigation.

Will Robots Undermine “Warrior Virtues”?

In many—if not most—societies, military service—and sometimes even war itself—is exalted because of its putative role in shaping character and developing particular virtues. War and military training provide an opportunity for the exercise and development of physical courage, fortitude, loyalty, willingness to sacrifice for the sake of the community, and even love and compassion. The idea that there is a connection between war and virtue also plays an important role in the culture of military organisations, which typically promote a conception of a professional warfighter who is supposed to cultivate and demonstrate “warrior virtues” (Olsthoorn 2005; Robinson 2007; Toner 2000; United States Department of the Navy 2007; Watson 1999). These virtues both serve the pursuit of victory in war and place limits on the ways in which wars are fought. “Good warriors” do not abandon their comrades in battle, nor do they massacre civilians (Silver 2006).

Taking humans “out of harm’s way” by allowing them to fight remotely will also remove them from the opportunity to exercise those virtues that are currently constitutive of being a good warrior. Most dramatically, it does not seem possible to demonstrate courage while operating a UMS from outside of the theatre of battle (Fitzsimonds and Mahnken 2007, pp. 100–101; Shachtman 2005). It is also hard to see how their operators could demonstrate a willingness to sacrifice themselves for the good of the unit. It will still be possible for operators to work hard and above and beyond the call of duty but it will not be possible for them to demonstrate any of those virtues that require the existence of serious negative consequences or physical threats as conditions of their possibility. More controversially, the limitations of UMS may make it extremely difficult for the operators to demonstrate love and compassion through their operations. Again, while they might go “above and beyond” to save the life of a particular comrade (or non-combatant) to such an extent that we might normally be inclined to attribute their actions to love or compassion, the geographical and psychological distance between the operator and the UMS means that these actions may nonetheless fail to demonstrate these virtues.

The loss of the opportunity for warfighters to demonstrate virtues may, of course, prove an acceptable price to pay to preserve their lives and limbs. Similarly, warfighters may be happy to live without a knowledge of what

it means to kill. What is less clear is whether military culture and military organisations can continue to function, in their current form at least, if it no longer becomes plausible to posit a connection between war and the character of those who fight it.⁶ The development of UMS may thus pose a serious threat to the self-conception of the military and perhaps even to the long term survival of the military services in their current forms.

Will Robot Weapons Make War Immoral?

The development of remotely operated weapon systems that would allow their operators to kill others without risk to themselves poses a serious challenge to the justification of war itself. Yale University Professor Paul Kahn has argued that unless it involves risk to both parties war is unethical (Kahn 2002). Where the asymmetry between opposing forces is large enough such that one party faces little or no risk, the moral exercise of force by the state possessing superior military might must take on the character of *policing*.

According to Kahn, the primary difference between these distinct modes of engagement concerns the categories of people who may legitimately be targeted with lethal force and the justification for doing so. In war it is legitimate to target and kill enemy combatants—those who are actively involved in the business of warfare (Walzer 2000, pp. 144–6). It is permissible to kill (some) non-combatants in the course of attacks on combatants as long as the non-combatants were not themselves directly targeted and reasonable care was taken to avoid and/or minimise non-combatant fatalities. In policing, on the other hand, it is permissible to kill only those who are responsible for serious violations of human rights and even then only in the course of preventing serious rights violations and/or trying to bring the perpetrators of such violations to justice. It is almost never morally permissible to kill persons who are not responsible for such rights violations in the course of policing operations (Kahn 2002).

The justification of the use of force in policing therefore proceeds by reference to the concepts of individual guilt and responsibility. The justification for killing in war is more complex and controversial. A full investigation of this topic is beyond the scope of this chapter. My ambition here is more limited. Following Kahn, I want to bring out the ways in which our intuitions about legitimate killing in war rely on the presupposition that opposing combatants are mutually at risk.

⁶ Fitzsimonds and Mahnken (2007) note that their enquiry, in the course of their survey of the attitudes of officers to UAVs, about whether the operators of UAVs should be eligible to receive the Distinguished Flying Cross for heroism in aerial flight “elicited a number of written comments ...” and that “One respondent termed the question itself ‘incredibly disturbing’” (p. 101), suggesting that contemporary operations of UAVs are already unsettling the self-conception of some military personnel.

To begin with, note that there are no morally relevant features that distinguish those it is permissible to directly kill (combatants) and those it is not (non-combatants) other than their membership in these groups. Combatants do not “deserve” to die any more than noncombatants. Many combatants do not support the war effort. Many non-combatants are fierce supporters of the war. Where conscription is practised, individuals may not even have chosen their status as combatants (Kahn 2002). Instead, the distinction between combatants and non-combatants is itself a product of war, which is a relationship between states (or state-like actors) (Sparrow 2005). It is war that makes combatants enemies and confers on them a licence to kill each other. Members of the armed forces of one state do not possess a licence to kill members of the armed forces of another state with which the state they serve is not at war.

Importantly, the fact that two states are at war does not confer on *all* their citizens the right to kill the “enemy.” It is only those in the service of the state in certain roles that lose their right not to be killed and who gain a licence to kill those of their enemies who serve in similar roles. Indeed, there is a sense in which non-combatants are not each others’ “enemies” at all (Lee 2004, p. 238). For instance, if an Australian and a Vietnamese civilian met during the American War in Vietnam and one killed the other, it would be an inadequate justification to explain that they were enemies because they were “at war.” What makes combatants enemies in a strong sense is that through their activities in the pursuit of the state’s military goals they are also working to kill each other (Murphy 1985). They stand in a relationship of “mutual threat” and their “licence to kill” stems from their right to self defence (Fullinwider 1985; Kahn 2002). Our ethical practices concerning the treatment of combatants who have been wounded or who have surrendered, in particular, strongly suggest that threat plays a crucial role in underpinning the right of combatants to kill other combatants. Surrender restores to combatants their right not to be killed. Similarly, because they have ceased to be a threat, severely wounded warfighters who have been overrun in battle are not legitimate targets in war.

However, if threat plays a central role in establishing warfighters “licence to kill” then, as Kahn argues, should the armed forces of one side of a conflict be rendered incapable of posing a threat to the other, then the warfighters of the more powerful side will no longer have a justification for targeting the armed forces of the other. Instead, the moral permissions and obligations of war are replaced by a more limited right to self defence alongside the permissions and demands of policing.

It is important to acknowledge that the relevant asymmetry is that which may exist between military forces considered as a whole. Kahn’s paper neglects this aspect of the argument and, by emphasising that combatants’ permission to kill each other is founded in their individual right to self defence, risks conveying the impression that it is the immediate threat to their person that is relevant. However, warfighters routinely kill others who pose no immediate threat to their persons in war, as, for instance, when they shell enemy positions over the horizon or bomb

munitions factories from a great height. Any plausible argument in support of their permission to do this cannot be founded on their right to defend themselves against *specific* threats. Instead, the threat that warfighters on opposing sides pose to each other and which justifies their licence to kill is a general one. *All* of the combatants on one side are a threat to *all* of those on the other. This feature of war is highlighted, in the course of satirising it, in the following section of dialogue from Joseph Heller’s *Catch 22*.

‘They’re trying to kill me,’ Yossarin told him calmly.

‘No one’s trying to kill you,’ Clevinger cried.

‘Then why are they shooting at me?’ Yossarian asked.

‘They’re shooting at *everyone*,’ Clevinger answered. ‘They’re trying to kill everyone.’

‘Who’s they?’ he [Clevinger] wanted to know. ‘Who, specifically, do you think is trying to murder you?’

‘Every one of them,’ Yossarian told him. (Heller 1994, p. 19)

However, that “they” are “trying” to kill me is not enough to justify my killing them in self defence; they must also have a realistic chance of doing so, such that they constitute a threat to my safety. If the military asymmetry between two sides to a conflict becomes large enough then the forces of the weaker side may not pose any such threat.

The presence of robotic weapons in a nation’s armed forces is neither necessary nor sufficient to establish asymmetry of this sort. Indeed, Kahn makes his argument without specific reference to UMS. What matters is the extent and nature of the asymmetry, not how it is achieved. However, by opening up the possibility that all human beings could be removed from the direct line of sight of the enemy and that indirect fire threats could be destroyed as soon as they appear in the battlespace, robotic weapons make it much more likely that the required asymmetry might be established. The only other foreseeable circumstances in which the armed forces of one side might pose no significant threat to the other might be a military campaign conducted entirely from the air.⁷ In either circumstance, if the argument that I am exploring here is to have any application it must be the case that there is some threshold of risk that establishes that the asymmetry between contesting forces justifies only policing and not war rather than the *complete* absence of risk. It is hard to imagine any conflict in which enemy combatants pose *no risk at all* to friendly forces. However, it *is* possible to imagine a circumstance in which the forces of one side presents so little threat to those of the other that warfighters on the superior side are not justified in directing lethal force against the warfighters of a conscript army merely because they are bearing arms.

⁷ This is the context in which Kahn (2002) developed his argument, discussing the NATO campaign over Bosnia. See also Kahn (1999).

It is possible, then, that future advances in the development of UMS may have the perverse consequence of undermining the justification of their use in war (or at least, war against certain sorts of enemies). Instead, robotic weapons would need to be used (much) more selectively, to prevent injustice and punish crime in the context of international policing operations.

Robotic Wars and the Future of Ethics

I have not attempted here to reach any overall conclusion “for” or “against” the development or application of military robots. This would be a futile task given the wide range of capacities of these systems and roles in which they might be used. Rather, my aim has been to draw attention to various implications of the development of robotic weapons for the future of war that I believe are significant for the ethics of war.

An important observation that emerges from this discussion is that it is far from obvious that all of the implications of the use of military robots make war “worse.” The argument that they will make killing more likely by allowing “killing at a distance” is less powerful than one might have thought. Moreover, insofar as these weapons allow warfighters an increased ability to discriminate between legitimate and illegitimate targets, they may both reduce the number of noncombatant casualties and also increase the demandingness of the principles of jus in bello. However, these possibilities need to be weighed against the possibility that the availability of UMS will make the use of force more likely, both at a tactical and strategic level (Sparrow 2009a). Ultimately, the impact of robotic weapons on how wars are fought is, of course, an empirical matter. Yet, given that controversy has erupted—and may be expected to continue to embroil—the introduction of these weapons, it is worth observing that the matter is more open than first appears.

I have also highlighted a number of other issues that are worthy of further investigation, both empirical and philosophical. In particular, there are reasons to expect the introduction of robotic weapons to impact on military culture in ways that might be extremely significant for the extent to which they encourage—or even allow—ethical behaviour. This is therefore an important topic for further study by those with skills to investigate it. The introduction of robotic weapons also raises difficult philosophical questions about the proper way of understanding the principles of jus in bello as applied to robotic—and especially autonomous—weapon systems and the ethics of “riskless warfare.” I hope that my drawing attention to them in this context will allow other policymakers, ethicists, and philosophers to progress these important debates in order to keep pace with the ethical challenges that are likely to arise as we move into the era of “robotic wars.”

References

- Adams, T.K. (2001), “Future Warfare and the Decline of Human Decision-Making,” *Parameters: US Army War College Quarterly* (Winter, 2001–2): 57–71.
- Arkin, R.C. (2009), *Governing Lethal Behaviour in Autonomous Robots* (Boca Raton: CRE Press).
- Asaro, P. (2008), “How Just Could a Robot War Be?” in P. Brey, A. Briggles and K. Waelbers (eds), *Current Issues in Computing and Philosophy* (The Netherlands: IOS Press), 50–64.
- Bender, B. (2005), “Attacking Iraq, from a Nev. Computer,” *Boston Globe*, April 3, A6.
- Borenstein, J. (2008), “The Ethics of Autonomous Military Robots,” *Studies in Ethics, Law, and Technology*, 2:1: Article 2. DOI: 10.2202/1941-6008.1036. Available at: <http://www.bepress.com/selt/vol2/iss1/art2>
- Braybrook, R. (2007), “Drones: Complete Guide,” *Armada International*, 31:3: 1–36.
- Cummings, M.L. (2004), “Creating Moral Buffers in Weapon Control Interface Design,” *IEEE Technology and Society Magazine* (Fall): 2–33, 41.
- Dunlap, Jr., C.J. (1999), “Technology: Recomplicating Moral Life for the Nation’s Defenders,” *Parameters: US Army War College Quarterly* (Autumn): 24–53.
- Excell, J. (2007), “Unmanned Aircraft: Out of the Shadows,” *The Engineer*, January 18.
- Featherstone, S. (2007), “The Coming Robot Army,” *Harper’s Magazine*: 43–52.
- Fitzsimonds, J.R., and Mahnken, T.G. (2007), “Military Officer Attitudes Toward UAV Adoption: Exploring Institutional Impediments to Innovation,” *Joint Force Quarterly*, 46: 96–103.
- Fullinwider, R.K. (1985), “War and Innocence,” in *International Ethics*, edited by C.R. Beitz, M. Cohen, T. Scanlon and A.J. Simmons (Princeton, N.J.: Princeton University Press).
- Graham, S. (2006), “America’s Robot Army,” *New Statesman*, 135:4796: 12–15.
- Hambling, D. (2007), “Military Builds Robotic Insects,” *Wired Magazine*, 23 January. Available at <http://www.wired.com/science/discoveries/news/2007/01/72543>
- Hanley, C.J. (2007), “Robot-Aircraft Attack Squadron Bound for Iraq,” *Aviation.com*, 16 July.
- Heller, J. (1994), *Catch 22* (London: Vintage).
- Hockmuth, C.M. (2007), “UAVs—The Next Generation,” *Air Force Magazine*, February, 70–4.
- Kahn, P.W. (2002), “The Paradox of Riskless Warfare,” *Philosophy & Public Policy Quarterly*, 22:3: 2–8.
- Kahn, P.W. (1999), “War and Sacrifice in Kosovo,” *Philosophy & Public Policy*, 19 :2/3.
- Kenyon, H.S. (2006), “Israel Deploys Robot Guardians,” *Signal*, 60:7: 41–4.

- Krishnan, A. (2009), *Killer Robots: Legality and Ethicality of Autonomous Weapons* (Farnham: Ashgate).
- Lee, S. (2004), "Double Effect, Double Intention, and Asymmetric Warfare," *Journal of Military Ethics*, 3:3: 233–51.
- Mander, J. (1978), *Four Arguments for the Elimination of Television* (New York: Morrow Quill Paperbacks).
- Meilinger, P.S. (2001), "Precision Aerospace Power, Discrimination, and the Future of War," *Aerospace Power Journal*, 15:3: 12–20.
- Metz, S. (2000), "The Next Twist of the RMA," *Parameters: US Army War College Quarterly*, Autumn: 40–53.
- Murphy, J.G. (1985), "The Killing of the Innocent," in J. Sterba (ed.) *The Ethics of War and Nuclear Deterrence* (Belmont, California: Wadsworth).
- Nagel, T. (1972), "War and Massacre," *Philosophy and Public Affairs*, 1: 123–44.
- Office of the Secretary of Defense (2005), *Unmanned Aircraft Systems Roadmap: 2005–2030* (Washington DC: Department of Defense, United States Government).
- Olsthoorn, P. (2005), "Honor as a Motive for Making Sacrifices," *Journal of Military Ethics*, 4:3: 183–97.
- Peterson, G.I. (2005), "Unmanned Vehicles: Changing the Way to Look at the Battlespace," *Naval Forces*, 26:4: 29–38.
- Robinson, P. (2007), "Ethics Training and Development in the Military," *Parameters: US Army War College Quarterly* (Spring): 22–36.
- Scarborough, R. (2005), "Special Report: Unmanned Warfare," *Washington Times*, May 8.
- Schmitt, M.N. (2005), "Precision Attack and International Humanitarian Law," *International Review of the Red Cross*, 87:859: 445–66.
- Shachtman, N. (2005a), "Attack of the Drones," *Wired Magazine* 13:6, Available at http://www.wired.com/wired/archive//13.06/drones_pr.html, at 25.08.05.
- Sharkey, N. (2009), "Death Strikes from the Sky: The Calculus of Proportionality," *IEEE Technology and Society*, 28:1: 16–19.
- Silver, S.M. (2006), "Ethics and Combat: Thoughts for Small Unit Leaders," *Marine Corps Gazette*, 90:11, 76–8.
- Singer, P.W. (2009), *Wired for War: The Robotics Revolution and Conflict in the 21st Century* (New York: Penguin Books).
- Sparrow, R. (2009a), "Predators or Plowshares? Arms Control of Robotic Weapons," *IEEE Technology and Society*, 28:1: 25–9.
- Sparrow, R. (2009b), "Building a Better WarBot: Ethical Issues in the Design of Unmanned Systems for Military Applications," *Science and Engineering Ethics*, 15:2: 169–87.
- Sparrow, R. (2007a), "Killer Robots," *Journal of Applied Philosophy*, 24:1: 62–77.

- Sparrow, R. (2007b), "Revolutionary and Familiar, Inevitable and Precarious: Rhetorical Contradictions in Enthusiasm for Nanotechnology," *NanoEthics*, 1:1: 57–68.
- Sparrow, R. (2005), "'Hands Up Who Wants To Die?': Primoratz on Responsibility and Civilian Immunity in Wartime," *Ethical Theory and Moral Practice*, 8: 299–319.
- Sparrow, R. (2004), "The Turing Triage Test," *Ethics and Information Technology*, 6:4: 203–13.
- Sullivan, J.M. (2006), "Evolution or Revolution? The Rise of UAVs," *IEEE Technology and Society Magazine*, 25:3 (Fall): 43–9.
- Toner, J.H. (2000), *Morals Under the Gun: The Cardinal Virtues, Military Ethics, and American Society* (Lexington: The University Press of Kentucky).
- Tuttle, R. (2003), "Kill Chain Timeline Now Down To 'Single Digit Minutes'," *Aerospace Daily*, 206:58: 5.
- Ulin, D.L. (2005), "When Robots Do the Killing," *Los Angeles Times*, January 30.
- United States Department of the Navy (2007), *Department of the Navy Core Values Charter 2007* [cited 17.7.2007]. Available from <http://ethics.navy.mil/corevaluescharter.asp>.
- Walzer, M. (2000), *Just and Unjust Wars: A Moral Argument with Historical Illustrations*. 3rd edn (New York: Basic Books).
- Watson, B.C.S. (1999), "The Western Ethical Tradition and the Morality of the Warrior," *Armed Forces and Society*, 26:1: 55–72.