



Shooting to Kill: The Ethics of Police and Military Use of Lethal Force

Seumas Miller

Print publication date: 2016

Print ISBN-13: 9780190626136

Published to Oxford Scholarship Online: November 2016

DOI: 10.1093/acprof:oso/9780190626136.001.0001

Autonomous Weapons and Moral Responsibility

Seumas Miller

DOI:10.1093/acprof:oso/9780190626136.003.0011

Abstract and Keywords

In this chapter the question addressed is whether autonomous robotic weapons necessarily compromise the moral responsibility of human combatants and their leaders. In order to answer this question, a novel argument is developed: the moral ramification argument. The conclusion of this argument is that it is highly improbable that the moral *jus in bello* principles of military necessity, discrimination, and proportionality could ever be programmed into robots. The argument utilizes theoretical descriptions of the key notions of military necessity, and individual and collective moral responsibility. Crucially, it relies on the inability of computers to detect and respond to moral—as opposed to physical—properties.

Keywords: autonomous weapons, moral responsibility, collective moral responsibility, moral ramification argument, military necessity, proportionality, discrimination, robots, computers

SCIENCE FICTION MOVIES, such as the *Terminator* series, have accustomed us to images of armed computerized robots led by leader robots fighting wars against human combatants and their human leaders. Moreover, by virtue of developments in artificial intelligence, the robots have superior calculative and memory capacity—after all, they are computers. In addition, robots are utterly fearless in battle, since they don't have emotions and care nothing for life over death. Does the human race, then, face robopocalypse? The short answer is no. Computers, robotic or otherwise, are not minded agents, steadfast intentional stances toward them notwithstanding.¹ Rather, these images are fanciful anthropomorphisms of machines; and the military reality is quite different.

Nevertheless, the specter of robocalypse persists, especially in the context of new and emerging (so-called) autonomous robotic weaponry. Consider, for example, the Samsung stationary robot that functions as a sentry in the demilitarized zone between North and South Korea. Once programmed and activated, it has the capability to track, identify, and fire its machine guns at human targets without the further intervention of a human operator. Predator drones are used in Afghanistan and the tribal areas of Pakistan to kill suspected terrorists. While the ones currently in use are not autonomous weapons, they could be, given this capability, in which case, once programmed and activated, they could track, identify, and destroy human and other targets without the further intervention of **(p.272)** a human operator. Moreover, more advanced autonomous weapons systems, including robotic ones, are in the pipeline.

In this chapter I explore the implications of autonomous robotic weapons, and related military weaponry, for the individual and collective moral responsibility of human beings engaged in war. Do such weapons necessarily compromise the moral responsibility of human combatants and their leaders, and, if so, in what manner and to what extent? In order to answer these questions I rely on serviceable theoretical descriptions of the key notions of war (Chapter 6, section 6.1) and military necessity, and individual and collective moral responsibility (Chapter 5, sections 5.2 and 5.3). In respect of the questions arising for individual and collective moral responsibility in respect of autonomous robotic weaponry, I provide what I refer to as the *moral ramification argument*. The conclusion of this argument is that it is highly improbable that moral *jus in bello* principles of military necessity, discrimination, and proportionality could ever be programmed into robots.

10.1 War, Collective Action, and the Principles of Necessity and Proportionality

As discussed in Chapters 3 and 6, waging war is typically morally justified by recourse to some notion of collective self-defense, such as defense of the nation-state against the armed aggression of another nation-state or of a nonstate actor such as a terrorist group. This ultimate end of collective self-defense and, relatedly, winning the war is necessarily underspecified prior to its realization. For example, the United States did not know when it declared war on Japan as a result of the Japanese attack on Pearl Harbor that victory over Japan would ultimately result from dropping atomic bombs on Nagasaki and Hiroshima. Moreover, the ins and outs of the evolving route leading to victory is also necessarily underspecified prior to its actually being taken; after all, it largely turns on what the enemy does, including by way of response to one's own armed attacks. So war is quite unlike programming a destination into a robot-driven car with a detailed and fixed roadmap, or, for that matter, a flight path into a computer-controlled jet aircraft. Nor is it even like playing a game such as chess, albeit it is analogous in some ways. For unlike in war, in chess there is a single, definite, unchanging, and mutually known "theater of war" (the

chessboard), a resource base that cannot reproduce itself (the **(p.273)** chess pieces), a sharply defined set of rules and contexts of application, and a fixed, finite, and knowable (at least in principle) set of possible moves and countermoves.

As we have seen, the actual conduct of war is governed by moral principles (the so-called *jus in bello* of just war theory), notably the principles of (1) military necessity, (2) proportionality, and (3) discrimination.² As will become evident, these are quite unlike the sharply defined rules and contexts of application in chess.³ For the moment I note that these principles have to be applied in very different military contexts, such as conventional theaters of war and counterterrorism operations, and that, as I argue below, their application is *radically context dependent*—so the conditions in which they ought to be applied cannot be comprehensively specified in advance of those conditions coming into existence. Importantly, unlike in the case of law enforcement, these principles apply at the collective level, as opposed to merely at the individual level. So the context of any, or at least most, applications of these principles is multileveled.

As we have seen, there is the individual level of a one-to-one encounter between a combatant and an enemy combatant, a firefight involving multiple combatants on both sides, and a battle involving possibly thousands of combatants over an extended period of time. Nor is this the end of the matter, for, as we all know, any given battle is merely a phase element in the overall war. So there are further collective levels governed by, for example, the collective end of winning the war, as opposed to merely winning one of the battles. Perhaps winning the war is describable as a level-three joint action.

The point to be stressed now is that, as argued in Chapter 6, the principle of military necessity, in particular, but also the principles of proportionality and discrimination, apply at the various conceptually distinct collective levels (e.g., the level of a battle or ongoing war fought by a military organization), and not simply at the level of an individual combatant's lethal action considered as a discrete, self-contained action (e.g., the **(p.274)** necessity to kill an enemy combatant who will otherwise kill oneself). Accordingly, the context for the application of these moral principles is a multilevel (individual and collective end) context. In essence, the principle of military necessity ultimately pertains to the long-term, necessarily underspecified collective end of winning the war, which generates in turn a nested, dynamic series of medium- and short-term collective ends, such as winning particular battles or firefights. These short- and medium-term collective ends are means to the long-term collective end of winning the war, albeit means in need of further specification, adjustment, or even abandonment in light of the responses to them of the enemy armed forces.

Accordingly, the principle of military necessity is to be understood, first, in terms of short-, medium-, and long-term means and ends—that is, in diachronic terms. Something is necessary in this sense if, comparatively speaking, it is both an efficient and effective means to an end, and there is no obviously superior means available. If it is the *only* means then it is *strictly* necessary. However, this is frequently not the case, and so to this extent “necessity” is correspondingly less strict. Second, the strength of the necessity to deploy a given quantum of lethal military force in, say, the context of a battle turns in large part on the moral weight to be accorded to the winning of that battle in light of its likely contribution to the ultimate (necessarily underspecified) collective end of winning the war (and, of course, the somewhat indeterminate moral weight to be attached to the latter). In the case of a crucial battle in the context of a war of collective self-defense, the military necessity to deploy a large quantum of lethal military force might be both strong (there is much at stake) and strict (it is the only available means).

What of the principles of proportionality and discrimination? These principles are obviously also to be applied at all levels collective and individual levels: whether it is a brief one-combatant-to-one-enemy-combatant exchange of fire, a firefight involving multiple combatants on both sides, a battle, or the war as a whole that is under consideration, it is morally impermissible to intentionally kill innocent civilians, to put their lives at unnecessary risk, or to knowingly cause disproportionate large numbers of civilian deaths. Naturally, what is at stake at each of these different levels, including the quantum of lives, can vary greatly but this does not affect the applicability of the principles.

The principles of military necessity, discrimination, and proportionality are conceptually interdependent, so that one cannot be correctly applied without attending to the requirements of the others. Roughly speaking, **(p.275)** the principle of discrimination forbids intentional targeting of innocent civilians⁴ and foreseeably and avoidably putting their lives at unnecessary risk. The latter clause conceptually implicates the principle of military necessity, for a risk to civilians is unnecessary if the use of lethal military force that constitutes this risk is not militarily necessary. Thus the principles of military necessity and discrimination are conceptually interdependent. Moreover, as we saw above, both principles must be applied at all individual and collective levels. Since these levels are interconnected by virtue of nested collective ends, the application of the principle of discrimination may well be a complex matter necessarily involving taking into account (1) the risks to civilians at these various levels and (possibly) adjudicating between them; (2) military necessity at these various levels and (possibly) adjudicating between them; and (3) adjudicating between points 1 and 2. For example, pursuing tactic A (aerial bombing) to realize the collective end of winning a battle might lead to many more civilian casualties in this present battle than pursuing tactic B (taking and holding ground without aerial bombing). However, pursuing A might be a more efficient and effective

means of decisively winning the battle (because, say, of the much heavier enemy casualties inflicted prior to the enemy's retreat), and might, therefore, reduce the number of future civilian casualties in future battles joined in further pursuit of the collective end of winning war.

The principle of proportionately arises in contexts in which both the principle of military necessity and the principle of discrimination are applicable. Roughly speaking, the legal principle requires that that the quantum of (unintended) civilian deaths resulting from the deployment of lethal military force should not be disproportionate to the strategic value, and the corresponding moral weight, of the collective military ends to be realized by that deployment. However, morally, as opposed to legally, there are three cohorts of potential lost lives to consider: civilians, one's own combatants, and enemy combatants. Naturally, military goals often involve maximizing enemy casualties in order to degrade enemy military capacity. So the means/end equation is in part to calculate the likely loss of life among one's own combatants (the cost) relative to the desired goal **(p.276)** to be achieved (maximum enemy casualties). But this calculation of military necessity is already one involving the application of a principle of proportionality. As such, the principle of proportionality (in our moral sense) is logically interdependent with the principle of military necessity.⁵ Moreover, as we saw above, the principle of proportionality applies at both the individual and collective levels. So the application of the principle of proportionality is complex in the manner of the other two principles.

This combination of logical interdependence between the three *jus in bello* principles and their applicability at all interconnected individual and collective levels in the overall context of a just war waged in collective self-defense gives rise to the phenomenon I refer to as *moral ramification*, and to the associated need for complex decision making of the kind described above. In short, in general one cannot simply apply one of these principles in a discrete, self-contained context (proportionality, for example, given the likelihood of heavy civilian casualties in a firefight), without taking into account the other principles and other contexts at other levels (e.g., the military necessity to win the battle in which the firefight is an important constitutive element).

Finally, I note that the moral considerations that arise from collective military ends at the collective level often outweigh, or otherwise render irrelevant, the moral considerations that arise at the individual level. In this respect, the deployment of lethal force by the military in war is quite different from the use of lethal force by police in law enforcement (see Chapters 3, 4, and 6).

10.2 Autonomous Robotic Weaponry and Human Moral Responsibility

Autonomous weapons are weapons system that, once programed and activated by a human operator, can—and, if used, do in fact—identify, track, and deliver lethal force without further intervention by a human operator. By “programmed”

I mean, at least, that the individual target or type of target has been selected and programmed into the weapons system. By “activated” I mean, at least, that the process culminating in the already programmed weapon delivering lethal force has been initiated. **(p.277)** This weaponry includes weapons used in nontargeted killing, such as autonomous antiaircraft weapons systems used against multiple attacking aircraft, or, more futuristically, against swarm technology (e.g., multiple lethal miniature attack drones operating as a swarm so as to inhibit effective defensive measures); as well as ones used or, at least, capable of being used in targeted killing (e.g., a predator drone with face-recognition technology and no human operator to confirm a match).

We need to distinguish between so-called “human-in-the-loop,” “human-on-the-loop,” and “human-out-of-the-loop” weaponry. It is only human-out-of-the-loop weapons that are autonomous in the required sense. In the case of human-in-the-loop weapons, the final delivery of lethal force (e.g., by a predator drone), cannot be done without the decision to do so by the human operator. In the case of human-on-the-loop weapons, the final delivery of lethal force can be done without the decision to do so by the human operator; however, the human operator can override the weapon system’s triggering mechanism. In the case of human-out-of-the-loop weapons, the human operator cannot override the weapon system’s triggering mechanism, so once the weapon system is programmed and activated there is not, and cannot be, any further human intervention.

The lethal use of a human-in-the-loop weapon is a standard case of killing by a human combatant, and as such is presumably, at least in principle, morally permissible. Moreover, other things being equal, the combatant is morally responsible for the killing. The lethal use of a human-on-the-loop weapon is also, in principle, morally permissible. Moreover, the human operator is, perhaps jointly with others, morally responsible, at least in principle, for the use of lethal force and its foreseeable consequences. However, these two propositions concerning human-on-the-loop weaponry rely on the following assumptions:

1. The weapon system is programmed and activated by its human operator and either:
2. (a) on each and every occasion of use, the final delivery of lethal force can be overridden by the human operator, and (b) this operator has sufficient time and sufficient information to make a morally informed, reasonably reliable judgment about whether or not to deliver lethal force, *or*
3. (a) on each and every occasion of use, the final delivery of lethal force can be overridden by the human operator, and (b) there is no moral **(p. 278)** requirement for a morally informed, reasonably reliable judgment on each and every occasion of the final delivery of force.

A scenario illustrating 3b might be an anti-aircraft weapons system being used on a naval vessel under attack from a squadron of manned aircraft in a theater of war at sea, in which there are no civilians present.⁶

What of human-out-of-the-loop weapons?⁷ Consider the following scenario, which I contend is analogous to the use of human-out-of-the-loop weaponry: There is a villain who has trained his dogs to kill on his command, and an innocent victim on the run from the villain. The villain gives the scent of the victim to the killer dogs by way of an item of the victim's clothing, and then commands the dogs to kill. The dogs pursue the victim deep into the forest, so that the villain is now unable to intervene. The dogs then kill the victim. The villain is legally and morally responsible for murder. However, the killer dogs are not, although they may need to be destroyed on the grounds of the risk they pose to human life. Hence the villain is morally responsible for murdering the victim, notwithstanding the indirect nature of the causal chain from the villain to the dead victim; the chain is indirect, since it crucially depends on the dogs doing the actual physical killing. Moreover, the villain would also have been legally and morally responsible for the killing if the "scent" was generic and, therefore, carried by a whole class of potential victims, and if the dogs had killed one of these individuals. In this second version of the scenario, the villain does not intend to kill a uniquely identifiable individual,⁸ but rather one (or perhaps multiple) members of a class of individuals.⁹

(p.279) By analogy, human-out-of-the-loop weapons—called "killer-robots"—are not morally responsible for any killings they cause.¹⁰ Consider the case of a human-in-the-loop or human-on-the-loop weapon. Assume that the programmer/activator of the weapon and the operator of the weapon at the point of delivery are two different human agents. If so, then, other things being equal, they are jointly (that is, collectively) morally responsible for the killing done by the weapon (whether it be of a uniquely identified individual or an individual qua member of a class).¹¹ No one thinks the weapon is morally or other than causally responsible for the killing. Now assume this weapon is converted to a human-out-of-the-loop weapon by the human programmer-activator. Surely this human programmer-activator now has *full* individual moral responsibility for the killing, as the villain does in (both versions of) our killer dog scenario. To be sure, there is no human intervention in the causal process after programming and activation. But the weapon has not been magically transformed from an entity only with causal responsibility to one which now has moral or other than causal responsibility for the killing.

It might be argued that the analogy does not work because killer dogs are unlike killer robots in the relevant respects. Dogs are minded creatures, whereas computers are not; dogs have some degree of consciousness and can experience, for example, pain. However, this difference would not favor ascribing moral responsibility to computers rather than dogs; rather, if anything, the reverse is

true. Clearly, computers do not have consciousness, cannot experience pain or pleasure, do not care about anyone or anything (including themselves), and cannot recognize moral properties, such as courage, moral innocence, moral responsibility, sympathy, or justice. Therefore, they cannot act *for the sake of* moral ends or principles *understood as moral in character*, such as the principle of discrimination. Given the nonreducibility of moral concepts and properties **(p. 280)** to nonmoral ones and, specifically, physical ones,¹² at best computers can be programmed to comply with some *nonmoral proxy* for moral requirements. For example, “Do not intentionally kill morally innocent human beings” might be rendered as “Do not fire at bipeds if they are not carrying a weapon or they are not wearing a uniform of the following description . . .” I return to this issue below.

Notwithstanding the above, some have insisted that robots are minded agents; after all, it is argued, they can detect and respond to features of their environment, and in many cases they have impressive storage, retrieval, and calculative capacities. However, this argument relies essentially on two moves that should be resisted and are, in any case, highly controversial. First, rational human thought (notably rational decisions and judgments) is downgraded to the status of mere causally connected states or causal roles, for example via functionalist theories of mental states. Second, and simultaneously, the workings of computers are upgraded to the status of mental states, for example via the same functionalist theories of mental states. For reasons of space I cannot here pursue this issue further. Rather, I simply note that this simultaneous downgrade/upgrade faces prodigious problems when it comes to the ascription of (even nonmoral) autonomous agency. For one thing, autonomous agency involves the capacity for non-algorithmic inferential thinking, such as the generation of novel ideas. For another, computers do not have interests or desires, do not pursue ends in themselves, and cannot choose their own ends. At best they can select between different means to the ends programmed into them. Accordingly, they are not autonomous agents, even nonmoral ones. For this reason alone, robopocalypse is evidently an illusion. Robotic weapons are morally problematic, but not for the reason that they are autonomous agents in their own right.

Granted that “autonomous” human-out-of-the-loop weapons are not autonomous (morally or otherwise), it has nevertheless been argued that there is no in-principle reason why they should not be used. (Moreover, they are held to have certain advantages over human-in-the-loop and human-on-the-loop systems—for example, being machines, they are not subject to psychological fear and associated stress.¹³) A key claim on which **(p.281)** this argument is based is that moral principles, such as military necessity, proportionality, and discrimination, can be reduced to rules, and these rules can be programmed in to computers.¹⁴ However, I suggest that the phenomenon of moral ramification presents a critical, if not insurmountable, problem at this point. To recap this

phenomenon: the combination of logical interdependence between the three *jus in bello* principles and their applicability at all interconnected individual and collective levels gives rise to *moral ramification* and the associated need for complex decision making, such that one cannot simply apply one of these principles in a given conceptually discrete and self-contained context involving the use of lethal force without taking into account the other principles and other contexts at other levels.

Let us revisit what this might mean in practice. Appropriate applications of, say, the principle of military necessity involves reasonably reliable, morally informed, contextually dependent judgments at the various collective levels, as well as at the individual level, and at the various centers of individual and collective responsibility. However, given the nested character of the individual and collective ends in play, their necessarily underspecified content, and the need to be responsive to the actions, including countermeasures, of enemy combatants and their leaders, there is a constant interplay between the various collective and individual levels (e.g., strategic commanders at headquarters and combatants in a firefight), and across centers (e.g., different theaters of war). Further, the various applications of the principles of necessity, proportionality, and discrimination are logically interdependent; for example, the application of the principle of proportionality depends on considerations of military necessity, and vice versa. Accordingly, there is a need to adjudicate not only between the means to given ends, but also with respect to the moral weight to be accorded different competing ends at different levels. For example, the individual end to advance to assist a comrade-in-arms coming under heavy fire might compete with the collective end of one's platoon or company to **(p.282)** make a tactical retreat to avoid heavy losses. Again, the collective military end to win firefights might be facilitated by relatively permissive rules of engagement (ROE), but perhaps this end competes with the collective end to avoid large-scale casualties among civilians, and the latter end is facilitated by relatively restrictive ROE. Further, at the macro-collective level, the collective end of the military leadership to win an internecine war might compete with the collective end of the political leadership not to inflict losses of a magnitude that would undermine the prospects for a sustainable peace.

In the light of this, let us see what it implies for the project to reduce the three *jus in bello* principles to rules and program them into armed robots. First, each moral principle needs to be expressible in a sharply defined rule couched in nonmoral descriptive terms. Given the nonreducibility of the moral to the nonmoral (physical?), it is extremely doubtful that this can be done for moral principles, especially ones that are relatively vague and quite general in form, as are the ones in question. Moreover, even if it could be done, the principles are *logically interdependent*, and this would need somehow to be accommodated; logically independent rule specifications, for example, would not work. Second, many, if not most, of the uses of lethal force in question are *joint* actions, and

joint actions are not reducible to aggregations of individual actions (see Chapter 1). So the rules in question would need somehow to accommodate this; the mere aggregation of instructions for single actors, for example, would not suffice. Third, the sharply defined rules in question would presumably be applicable to sharply defined, discrete, self-contained contexts involving the use of lethal force; otherwise the robot would not be able to comply with them. Here the phenomenon of moral ramification comes fully into its own. For, as our above examples demonstrated, in any such conceptually discrete and self-contained context, be it a one-against-one encounter, a firefight, an air strike, or a battle, there will inevitably be moral considerations emanating from some other context (for example, another battle) or some larger context of which the discrete, self-contained context is an element (for example, the war as a whole), which bear upon it in a manner that morally overrides or qualifies compliance with the sharply defined rule in question (or set of rules, for that matter¹⁵). Given that each war (p.283) taken in its totality, is unique, this interplay of contexts has the effect of making decisions in accordance with the *jus in bello* highly, indeed, radically, contextually dependent, and particular in character. As such, these decisions are beyond the reach of rules, however sharply defined; for rules are necessarily general in character and defined prior to their application in particular contexts. I conclude that this “computerized” conception of the application of fundamental moral principles in war faces prodigious, if not insurmountable, problems. In short, evidently robopocalypse is doubly an illusion.

An important consequence of this is that the design, construction, and use of human-out-of-the-loop weapons are highly morally problematic. Such weapons cannot be programmed to comply with the moral principles of military necessity, discrimination, and proportionality. Moreover, their use would seriously impede the capacity of their human operators to adequately comply with these moral principles, and, to this extent, it would be an abnegation of moral responsibility on the part of the military. Finally, the use of these human-out-of-the-loop weapons is evidently unnecessary, since, as we saw above, for the combat situations in which human-in-the-loop weapons are inadequate, human-on-the-loop weapons are available, and I conclude that human-out-of-the-loop weapons morally ought not to be used.

10.3 Conclusion

In this chapter I have addressed the question of whether autonomous robotic weapons necessarily compromise the moral responsibility of human combatants and their leaders. In order to answer this question, I have developed a novel argument: the moral ramification argument. The conclusion of this argument is that it is highly improbable that moral *jus in bello* principles of military necessity, discrimination, and proportionality could ever be programmed into robots. The argument utilizes theoretical descriptions of the key notions of war and military necessity (Chapter 6), and individual and collective moral responsibility

(Chapters 2 and 5). Crucially, it relies on the inability of computers to detect and respond to moral properties.

Notes:

(1.) Daniel Dennett, *The Intentional Stance* (Cambridge, Mass.: MIT Press, 1987).

(2.) There are various different possible formulations of and complications arising from these moral principles. For example, I will be concerned with proportionality as it pertains to civilian deaths. See Miller, *Terrorism and Counter-Terrorism*.

(3.) This point is not vitiated by the fact that these ius in bello principles are further specified by the ROE. For the problem remains, both at the level of the 'derivation' of the ROE from the ius in bello principles and in the application of the ROE themselves. In order to avoid unnecessary complications in this chapter I discuss the issues purely in terms of the ius in bello principles.

(4.) Arguably, the component clause of the principle of discrimination—the impermissibility of intentionally killing innocent civilians—is logically independent of its second clause and of the other principles. This does not affect my argument. The principle of discrimination also applies to the kind of weaponry used. For example, biological weapons are indiscriminate.

(5.) I am not arguing that the principle of proportionality as it applies in standard cases of personal self-defense is logically interdependent with the principle of necessity as that principle applies in such cases.

(6.) There are various other possible such scenarios. Consider a scenario in which there is a single attacker on a single occasion in which there is insufficient time for a reasonably reliable, morally informed judgment. Such scenarios might include ones involving a kamikaze pilot or a suicide bomber. If autonomous weapons were to be morally permissible, the following conditions at least would need to be met: (1) prior clear-cut criteria for identification and delivery of lethal force to be designed into the weapon, and used only in narrowly circumscribed circumstances; (2) prior morally informed judgment regarding criteria and circumstances; and (3) ability of operator to override the system. Here there is also the implicit assumption that the weapon system can be "switched off," which is not the case with biological agents released by a bioweapon.

(7.) See Ronald Arkin, "The Case for Ethical Autonomy in Unmanned Systems," *Journal of Military Ethics* 9, no. 4 (2010): 332–341. Arkin has argued in favor of the use of such weapons.

(8.) It is not a targeted killing.

(9.) Further, the villain is legally and morally responsible for foreseeable but unintended killing done by the killer dogs in the forest, if they had happened upon one of the birdwatchers well-known to frequent the forest and mistakenly killed him instead of the intended victim. (Perhaps the birdwatcher carried the scent of birds often attacked by the killer dogs).

(10.) Rob Sparrow, "Killer Robots," *Journal of Applied Philosophy* 24, no. 1 (2007): 63-77. For criticisms, see Uwe Steinhoff, "Killing Them Safely: Extreme Asymmetry and Its Discontents," in Strawser, ed., *Killing By Remote Control* 179-209.

(11.) Moreover, each is fully morally responsible; not all cases of collective moral responsibility involve a distribution of the quantum (so to speak) of responsibility. See Miller, "Collective Moral Responsibility."

(12.) The physical properties in question would not only be detectable in the environment, but also be able to be subjected to various formal processes of quantification and so on.

(13.) Arkin, "The Case for Ethical Autonomy in Unmanned Systems."

(14.) Arkin, "The Case for Ethical Autonomy in Unmanned Systems." This claim has been countered by various critics, but not, in my view, decisively. For these critics have, as far as I am aware, relied on piecemeal objections (so to speak), such as the difficulty an autonomous weapon would have in distinguishing innocent civilians from terrorists in civilian dress. See, for example, Noel Sharkey, "Killing Made Easy: From Joysticks to Politics," in Patrick Lin, Keith Abney, and George A. Bekey, eds., *Robot Ethics: The Ethical and Social Implications of Robotics* (Cambridge, Mass.: MIT Press, 2012), 111-128. However, a more decisive, and by contrast, *holistic* objection can be made to the application of these principles: the moral ramification argument.

(15.) The sharply defined computerized rule conception could be complicated by adding meta-rules, for example. However, this would not make any material difference to the problems; it would simply elevate things to a higher level of complexity.

Access brought to you by: