

Article

Autonomous Weapons Systems and Force Short of War

Nathan Wood ¹¹ Ghent University; Nathan.Wood@UGent.be

Abstract: Though war is never a good thing, all things considered, there are times when it is arguably justified. Most obviously, providing direct military assistance to a victim of unjust aggression would constitute a rather clear case for military intervention. However, the providing of direct military assistance may in some cases be a prospect fraught with risks and dangers, rendering it politically (and possibly even morally) difficult for states to adequately justify such action. In this article I argue that autonomous weapons systems present a way past this dilemma, providing a method for delivering direct military assistance, but doing so in a way that is less politically overt and hostile than sending one's own combat units to aid a beleaguered state. Thus, sending autonomous weapon systems (AWS) presents an additional forceful measure short of war which states may employ, adding to the political options available for combating unjust aggression, and allowing one to provide direct assistance to victim states without necessarily bringing one's own state into the conflict. In making this argument I draw on the current Russian invasion of Ukraine as a running example.

Keywords: Autonomous Weapons Systems; Military Ethics; Force Short of War; Russo-Ukrainian War

Citation: Wood, Nathan. 2022.
Autonomous Weapons Systems and
Force Short of War. *Journal of Ethics
and Emerging Technologies* 32: 2.
<https://doi.org/10.55613/jeet.v32i2.115>
5

Received: 17/10/2022
Accepted: 30/11/2022
Published: 31/12/2022

Publisher's Note: IEET stays neutral
with regard to jurisdictional claims in
published maps and institutional
affiliations.



Copyright: © 2021 by the authors.
Submitted for possible open access
publication under the terms and
conditions of the Creative Commons
Attribution (CC BY) license
(<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Weapons of war are becoming increasingly autonomous, and though such advancements create opportunities for clear gains, both strategically and morally, from some groups there remains staunch opposition to the development of what are known as autonomous weapons systems, or AWS. The arguments against AWS are varied, from principled concerns that such weapons would violate the human dignity of those targeted by them to more pragmatic worries that reducing the bloodshed of war might lessen its spectre, making armed conflict more likely. In this article I argue that, far from making war more likely, AWS in fact provide an additional means by which states can exert force short of war, thereby securing strategic or moral goals while avoiding the need for a declaration of war (and all that would follow from that). In particular, there exist situations where states are forced (either by circumstances or an adversary's military doctrine) to either eschew all direct military engagement or wage an all-out war. However, AWS present a way between these alternatives, allowing states to effectively militarily intervene without actually committing any forces to a conflict. This is most relevant for situations where there exist strong political, strategic, or moral reasons against war, but where military engagement is, all things considered, the right course of action. The way AWS (can be designed to) function also allows states to deploy them in deescalatory ways, avoiding the need to start a war in the first place, while also simultaneously reducing the likelihood of another party starting war. Thus, AWS serve only to add to a state's political and military options, and do so in a way that opens avenues for securing peace and avoiding future conflict.

The arguments are structured as follows. I begin in Section 2 by briefly sketching the objection that AWS will make wars more likely, and show that this objection is susceptible to a variety of rebuttals. In Section 3 I then present and defend the position that AWS can in fact be utilized as an alternative to war, providing one further argument in their favor and

one further response to the objection dealt with in the preceding section. Building on these points, in Section 4 I demonstrate how AWS can be used as a response to escalatory actions by aggressive states, and how the deployment and diffusion of AWS to allied states can serve as a persistent tool for deescalation. The arguments for AWS as a peacekeeping alternative to war (and a response to aggression) are then given their most forceful defense in Section 5, where I tackle a persistent objection to AWS and show that even the most morally hazardous form of AWS can be used in a manner that accords with the laws of armed conflict (LOAC) and furthers peace. It is also worth noting that throughout the arguments to come, the current Russian invasion of Ukraine is used as a running example to illustrate a number of arguments being developed.

Before moving onto these arguments, however, it is important to dispense with one definitional point. In the literature on AWS “multiple definitions and understandings currently exist about autonomous systems”,¹ and without making clear at the outset what definition one is using, it is possible (if not likely) that misunderstandings will follow.² Now, the competing positions one can find from scholars are by and large rooted in differing accounts of autonomy, or debates about the relation between “autonomy” and “autonomous”,³ “automated” and “autonomous”,⁴ or other related concepts.⁵ However, for all of the disagreement amongst scholars, there is a growing consensus between governments and non-governmental organizations (NGOs) concerning what definition of AWS is most apt. These organizations have by and large accepted a broad understanding of autonomous weapon systems, following the general position taken by the United States Department of Defense (DoD), which defines AWS as

weapon system[s] that, once activated, can select and engage targets without further intervention by a human operator. This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation.⁶

In a 2014 expert meeting of the International Committee of the Red Cross (ICRC), a similar view was expressed, maintaining that

‘autonomous weapon systems’ were defined as weapons that can independently select and attack targets. These are weapon systems with autonomy in the ‘critical functions’ of acquiring, tracking, selecting and attacking targets.⁷

In May 2021 the ICRC further clarified its view of AWS with the following definition:

Autonomous weapon systems select and apply force to targets without human intervention. After initial activation or launch by a person, an autonomous weapon system self-initiates or triggers a strike in response to information from the environment received through sensors and on the basis of a generalized

¹ Williams (2015) p. 27.

² See Wood *forthcoming* for elaboration of this point.

³ Roff (2013) pp. 353--354; Roff *forthcoming*.

⁴ Boothby (2016) pp. 247--257.

⁵ Williams (2015). See also Sec. III of Dinstein and Dahl (2020) for discussion of how the LOAC applies (or should apply) to AWS.

⁶ US Department of Defense (2017) pp. 13--14.

⁷ International Committee of the Red Cross (2014) p. 5.

“target profile”. This means that the user does not choose, or even know, the specific target(s) and the precise timing and/or location of the resulting application(s) of force.⁸

These definitions have many points which could be given nuanced treatment,⁹ but the core element is that AWS are understood as weapon systems that “can select and engage targets without further intervention”. Since this essential view captures the understanding of both the United States government and the single most important body of regulation for the laws of war, namely the ICRC, it is this definition with which we will move forward.¹⁰

2. AWS and the Likelihood of War

As mentioned above, one worry surrounding AWS is that by reducing the bloodshed of war, particularly the harm one’s own side may suffer, wars may become more likely,¹¹ a fact which proponents of this objection take to be a morally bad thing. However, there are solid grounds for resisting this objection, and on each and every of its underlying premises.¹²

First of all, though casualties among one’s own troops present a strong deterrent to war (at least in most rights-respecting states), there are other factors which overdetermine against war already. This is evidenced by the simple fact that while wars do occur more often than we might wish, they still occur far less often than would be the case if every country which could wage a casualtyless war did so whenever they could. As an example, the United States could surely wage casualtyless wars of conquest against the Bahamas, Barbados, or Grenada, and most likely against Jamaica, Haiti, and the Dominican Republic as well. The simple fact is that the overwhelming air and naval superiority the U.S. enjoys in the region means it could shell and bomb almost any Caribbean state with impunity until they inevitably surrendered. The reason the United States does not go to war with these countries has nothing to do with whether or not these countries can inflict casualties on American forces. And the same fact obtains all around the world; it is a credit to human civilization that we have made the economic and political costs of war so potentially high, and that we appreciate the extreme moral costs of war so thoroughly, that these alone can eliminate the possibility for wars, even wars that would be certain victories with almost no military costs (comparatively speaking) for potential aggressors.

One might rejoin that not all states are as effectively reigned in by political or economic concerns as generally rights-respecting regimes are. In point of fact, some states will carry out military adventures which are certain to bring grave losses in political capital and short- and medium-term economic losses as well, if they think their military chances are secure enough. As an example, accounts from both before and after the

⁸ International Committee of the Red Cross (2021) p. 1.

⁹ See Williams (2015) for just such an in-depth discussion.

¹⁰ Note that though this definition enjoys a growing number of adherents in international relations and military circles, it is still not universally accepted. For example, the U.K. Ministry of Defence defines AWS as weapon systems which are “capable of understanding higher-level intent and direction”, a much stricter view of AWS which does not agree with the broader conception laid out by the U.S. DoD or the ICRC. This, however, only highlights the need for definitional clarity at the outset. See UK Ministry of Defence (2017) p. 13, for their full definition.

¹¹ See, e.g., Asaro (2008) pp. 56--58; Gebrud (2014) p. 39; Heyns (2013) pp. 11--12; Sharkey (2012) p. 122; and Docherty et al. (2018) p. 6.

¹² The points developed below consist of a brief recapitulation of those made by Maciej Zająć, and by no means constitute the whole range of responses one may give to this objection. For Zająć’s statements, his wider consideration of examples, and his expanded arguments, see Zająć (2022a) Ch. 9.1.

February 24th invasion of Ukraine at least partly attribute Russia's decision to invade on a feeling of war optimism, a "sense that the [Ukrainian] government would fall with just a little push".¹³ Access to advanced AWS may therefore only heighten the chances of war breaking out between such regimes and those they may wish to conquer, or so the objection goes.

This rebuttal appears to have some bite; if the economic and political costs of war will alone not restrain aggression on the part of Russia (or similarly-minded states), then reducing the casualties of war might just tip the balance in favor of such a decision when this would otherwise not be the case. Thus, the ability of AWS to reduce the bloodshed of war might increase war's likelihood. However, one should not overstate this either. Russia may not be sufficiently moved by political or economic consequences, but they are also clearly not sufficiently moved by friendly losses either. If they were, the war in Ukraine would be over already, for Russia has seen more than enough men die there for it have halted their efforts if casualties alone had any impact on that state's military decision-making. The sad fact is that those states unmoved by political, economic, or moral arguments against war are also usually unmoved by the deaths of their own soldiers.

The persistent objector could rejoin that AWS still add a potential source of firepower which might be hard to counter, and this increased tactical advantage may make aggressive states more likely to enter wars, simply because their chances look better. Now, this may be true, but it is true for any military advantage one gains; if AWS should be banned because they present a means to win more easily or with a greater chance of success, then so should artillery be banned, and aircraft, and armor, and every weapon which presents an advantage. In a related vein, if the initial objection is granted, that reduced bloodshed makes wars more likely, then this would also argue against other technologies or actions that might reduce one's own casualties; field medical units, rapid removal of the wounded, body armor, or any other form of protective equipment lessens the spectre of war, and thus all should be eschewed.¹⁴ One might bite these bullets, but to do so is not just to argue against AWS, but against virtually every aspect of modern warfare. There may be intellectual merit in exploring such ideas, but this goes far beyond the scope of debates about AWS development and potential deployment.

As a final point before moving on, the initial objection above was that AWS could make wars more likely, *and this is a bad thing*. The objection thus crucially depends on the premise that morally speaking, the less wars, the better. However, this premise is not so clearly tenable as it might at first seem. While actions like Russia's invasion of Ukraine (in 2014 or 2022) or their earlier invasion of Georgia are clearly bad things, not all wars are wars of aggression and conquest. Moreover, those wars which are most morally justified are usually those of a humanitarian nature, where the states fighting to uphold rights have little to gain themselves. Because of this lack of gain (or gain of only a purely cosmopolitan moral value), such humanitarian wars are often hard to justify to a public back home, especially when they see their sons and daughters dying for the sake of unknown people half a world away. Yet failing to fight these wars can be far more devastating than certain unjust conflicts which do occur. As an example, there were only six deaths resulting from the 2014 Russian seizure of Crimea, an operation lasting just over one month,¹⁵ and from April 2014 through the end of December 2021 the United Nations Office for the High Commissioner for Human Rights estimated that there were roughly 14,200--14,400 deaths as a result of the Russian-backed separatist conflict in eastern Ukraine.¹⁶ Taken all

¹³ Kirby and Guyer (2022). Michael Kofman, an expert on Russia's armed forces, expressed similar conclusions on the eve of the invasion. See Kofman (2022).

¹⁴ This latter argument is forcefully made by Strawser (2010).

¹⁵ Kofman et al. (2017) p. 12.

¹⁶ United Nations Office of the High Commissioner for Human Rights (2022) p. 3.

together, this means that prior to the 2022 invasion of Ukraine, the Russo-Ukrainian war claimed an average of 154 lives per month.¹⁷ In comparison, the Rwandan Genocide saw between 500,000 and 800,000 civilians murdered in a roughly 100-day period.¹⁸ Taking the lowest estimate for the genocide, that means that roughly 5,000 civilians were murdered *per day* in an internecine conflict using the most rudimentary of weapons, and where virtually any outside military force could have been decisive (or at least created safe corridors for fleeing Tutsis). In fact, speaking nineteen years after the genocide, former president Bill Clinton lamented that if the U.S. and U.N. had acted, he believes they could have saved “at least a third of the lives that were lost... 300,000 of those people's lives”.¹⁹ Yet neither the U.S. nor U.N. did act, and for the simple reason that only a year prior nineteen U.S. servicemen had been killed during the UNOSOM II peacekeeping mission in Somalia, leaving those “at peacekeeping headquarters... [with] no major appetite to get involved in such missions”.²⁰ To reiterate, because of the deaths of nineteen men there was not enough political will at home to save half a million innocent civilians. And as Halabja, Srebrenica, Chechnya, Darfur, Sri Lanka, Myanmar, and South Sudan (among others) show, genocide and genocidal killings occur far more often than we would like, and usually with little outside response due to the political hurdles involved.

The deployment of AWS would not require that same political will, and so that hurdle would not exist. Thus, AWS can provide a politically viable means to aid would-be humanitarian interveners in actually carrying out the missions (and wars) they know they should. Thus, AWS may make *certain types of wars* more likely, namely wars of humanitarian intervention which protect innocent lives but provide no benefit to the interveners. This, however, hardly seems to provide a meaningful objection to AWS.²¹

3. Alternatives to War

So AWS might make *some wars* more likely, but these will almost exclusively be humanitarian interventions where a higher likelihood of participation is actually a good thing. Yet AWS can also reduce the likelihood of war by providing one more alternative to war available to those who wish to protect rights or aid states in defending against unjust aggression. This is because AWS are weapons,²² and as weapons they may be deployed by our military *or* be sold, traded, or loaned to other militaries.

This is not just a geopolitically important aspect of AWS, but a morally important one as well. Though many (if not most) wars have historically messy causes leading up to

¹⁷ Given the ongoing nature of the current engagements in Ukraine, and the extraordinary difficulty of collecting reliable information in such an environment, we will not attempt to examine the casualty figures arising since the February 24th invasion. Moreover, the UNOHCHR estimates for eastern Ukraine are likely to not be fully accurate either, due to limitations in access for outside observers.

¹⁸ Guichaoua (2020); Meierhenrich (2020).

¹⁹ CNBC Meets (2014).

²⁰ Interview with Philip Gourevitch (PBS Frontline (1999)). See also Barnett (2002) for a similar assessment.

²¹ To reiterate, the preceding points all see much greater development in Zajac (2022a) Ch. 9.1, which the interested reader is encouraged to explore.

²² Some definitions of AWS consider them to be something more than mere weapons (e.g., Sparrow (2007); Anderson and Waxman (2012); Roff (2013)). However, the types of systems we have now and those of the near future are nowhere near sophisticated enough to be treated as anything other than as weapons (Sharkey (2010)). Moreover, the U.S. DoD and the ICRC definitions are broad enough that whole classes of what are obviously weapons or ordnance fall under the definition of AWS. See, for example, Boulanin et al. (2020); International Committee of the Red Cross (2021a); or Wood *forthcoming*.

them, rendering it difficult to tell who exactly is in the wrong, there are sometimes situations where there is an obvious aggressor or grave abuser of rights. In these cases, those who can help arguably have a responsibility to do so.²³ However, even when states are under some obligation to help militarily and in fact have the general political will to do so, there may still exist powerful countervailing reasons which can prevent them from doing their best. For example, states under a *prima facie* obligation to militarily assist those resisting aggression may find that such direct military assistance is almost certain to escalate the conflict, causing it to be more bloody and deadly than if they were to do nothing at all.²⁴ In that case, they arguably ought not to send their own troops and risk escalation, even if that would be the surest way to halt and repel the aggressor forces. However, sending AWS to the attacked state, especially highly advanced AWS, provides a potential means for doing nearly as much as if we sent our own troops, but without the attendant high risks of escalation and subsequent political fallout.

To more clearly see this potentiality, let us return to the case of the ongoing Russo-Ukrainian war. In the immediate aftermath of Russia's invasion of Ukraine, there was across the world a clear and loud statement that this was beyond the pale and should be stopped immediately,²⁵ a sentiment doubly voiced by western and in particular NATO states, who levied astoundingly stark sanctions against Russia.²⁶ These economic strikes were not the only response either, as a flow of intelligence and arms to the Ukrainian military showed the commitment of NATO states and their allies to stopping Russia's aggression. However, despite the material and advisory aid that was and continues to be sent, no outside power has joined Ukraine in its fight, and this is most likely due to a fear of escalation. As one analyst put it, "NATO is currently unable to use force, even for humanitarian purposes, for fear of an escalatory response from Vladimir Putin".²⁷ The essential problem is that while we know our forces could (and most assuredly would) have been able to halt the aggression early on in the conflict, committing our troops to such an operation would have pitted NATO fighting units against those of Russia. In point of fact, NATO would have entered into war against Russia. The extraordinary risks of escalation in such a case should be evident, as well as why NATO has ultimately not committed to the fight so actively.

That risk of escalation, however, only attends the prospect of sending *combat units*. Importantly, as we have seen, the provisioning of advanced armaments and heavy weapons is something which can be done with minimal risks of escalation, and which does greatly improve the outlook for Ukrainian armed forces. Moreover, advanced systems which act as force multipliers or which cover tactical needs Ukraine cannot effectively secure for itself have allowed the invaded state to not just hold its defenses, but indeed begin retaking land occupied by Russia (or its puppets) for over eight years. AWS add one more type of such systems, and one which can require fewer operators in order to be used effectively, something crucial in a war where the aggressor has the potential to

²³ See Cooper and Kohler (2009) for a general introduction to the international norm of a "responsibility to protect". Evans (2009) provides more detailed examination of the principle itself.

²⁴ Note that this is not to say that the alternatives or comparisons for proportionality purposes are between "military intervention" and "doing nothing" (see, e.g., Wood's unpublished manuscript, "Rethinking Proportionality"). This is only to illustrate that sometimes intervention, despite its potential justness, can lead to consequences far worse than if one left aggression unchecked. Such would, for example, be the case if military intervention were certain to lead to a nuclear exchange.

²⁵ UN News (2022).

²⁶ Toh et al. (2022).

²⁷ Raine (2022).

field far more units than the victim could.²⁸ More than this, AWS designed for point-defense uses or to provide area denial to aircraft, armor, or naval vessels could allow Ukraine to focus its human forces and capitalize on tactical gains in one area while leaving other fronts to be held by autonomous systems. Such moves would provide strategic and tactical options to Ukraine, and allow for AWS to be deployed by them in situations where there is a lessened risk of the autonomous systems acting in a problematic fashion. I.e., if Ukraine were to leave anti-air and anti-armor AWS to hold a certain region against those forms of enemy systems, these AWS could operate with virtually no risk that civilians would be targeted. And importantly, it is worth stressing again that states providing AWS would be providing nothing more than advanced armaments, not troops. This would therefore make such provisioning a far less escalatory action than the sending of combat units would be.

Now, one might object that advanced AWS are apt to be complicated systems, and cannot simply be sent with a simple “how-to” manual. More specifically, the objection is that some, if not many, of the most effective forms of AWS are likely to require trained operators, thus making it impossible to send them without some form of crew. Therefore, AWS may not present such a straightforward means to immediately aid a victim state while avoiding potentially escalatory deployments of forces. This is a nuanced objection which pays heed to the realities of warfare; AWS are not bolt-action rifles that can be pushed into someone’s hands and expected to be used effectively. Moreover, more sophisticated AWS with more open targeting parameters are likely to require trained operators in order for their use to even be legal. However, the objection is not properly one against AWS in general, but only against *particular types* of AWS.

Recall that the definition of AWS we are using, namely that of the U.S. DoD and the ICRC, is a broad one. Under that broad definition, science fiction *Terminators* will count as AWS, but so will many types of homing munitions, so-called “fire-and-forget” systems, or point-defense batteries. The former sort of things – *Terminators* or any other systems which can be given broad goals or relative freedom in operation – are definitely apt to need very competent handlers who can ensure they are functioning within the LOAC. However, homing munitions, “fire-and-forget” missiles, and point-defense batteries can all be made to be AWS, and none need be greatly sophisticated or unfamiliar to militaries potentially receiving them. For example, two weapon systems which have had a significant impact on Ukraine’s self-defensive efforts are the Next Generation Light Anti-Tank Weapon (NLAW) supplied by Sweden and Great Britain and the Bayraktar TB2 unmanned aerial vehicle (UAV) built by Turkish defense company Baykar. Neither of these systems would satisfy the criteria for being deemed fully autonomous weapon systems,²⁹ but both could rather straightforwardly be modified to allow for full autonomy (in the DoD/ICRC understanding). For the case of the NLAW, this could be achieved by making it so operators only had to fire in the rough direction of enemy armor positions, after which the weapon itself would select and engage potential targets based on its own sensors, rather than requiring that a lock on enemy vehicles be first obtained by a human user. For the TB2, it would suffice to add simple tracking and targeting software which allows operators to select a few types of vehicles to be targeted. When the TB2’s sensors identify such vehicles, perhaps with an added precaution that it must also identify some sort of enemy marking like a combat insignia or enemy colors, then it would engage. These would then both be truly autonomous weapon systems (under the DoD/ICRC definition),

²⁸ Russia has roughly five times more active personnel than Ukraine. See International Institute for Strategic Studies (2022).

²⁹ The TB2 is fully remotely controlled by a ground crew, and the NLAW would at most count as a semi-autonomous system, given that the operator must first manually select a target and allow the weapon to gain a lock, after which the NLAW can self-guide its payload to the target. See US Department of Defense (2017) p. 14.

but neither would necessarily demand that instructors or foreign operators be present for these to be used effectively or legally. In fact, one of the great advantages of the NLAW is that it takes less than a day of training to be able to use it,³⁰ and the Ukrainian military is already familiar with the use of the TB2. Ostensibly, AWS variants of these systems would not greatly differ from the current versions they have, and so both of these could be sent to Ukraine without necessitating the deployment of any troops or advisors.

One might rejoin that even if AWS can ultimately be used in this fashion, especially more rudimentary AWS, then while this might work, it would only work once. Put differently, we can highlight that AWS are just another type of weapon system, and with that in mind, send them on the understanding that this is a politically viable and non-escalatory alternative to war. Doing so, we might even be able to help Ukraine now and do so in a way that is genuinely not escalatory. However, in response to such actions, Russia (or any other state, for that matter) may just make it a point of military doctrine that the provisioning of AWS to another state is considered equivalent to the sending of one's own forces. In that case, any future selling, trading, or loaning of AWS to beleaguered states would be considered an act of war, making it again escalatory and no longer a proper alternative to war. Very briefly, one can easily reply that even if it only works once, it still works once, and that makes it worthwhile. However, even granting the underlying idea of this objection, it too is overstating its case. It is entirely plausible that some states might consider highly advanced AWS as so akin to flesh-and-blood combat units that they make it a point of military doctrine that they would be treated as equivalent (and perhaps legitimately so). However, as with the preceding objection, this only really points to a potential issue with *highly advanced* autonomous systems. Yet virtually all currently existing AWS and much of those to be seen in the coming decades are by and large advanced autonomous munitions, point-defense turrets, or simple area denial platforms. These systems are not ones which can plausibly be treated as akin to combat units, nor are they systems which potential aggressors will necessarily even want to be treated as akin to combat units, as this would limit their ability to deploy such things as well. Moreover, the entire discussion regarding how advanced an AWS is points to a general problem with this objection, namely that one cannot treat the "provisioning of AWS" as the same as sending combat units or as an act of war, for the simple fact that there are so many types of systems which might count as AWS. Many properly autonomous systems will moreover be rather simple and routinely sold, traded, or loaned to other militaries (like, for example, advanced autonomous munitions). No country wants it to be the case that selling arms automatically entangles them in wars which they have no intention to join, but a doctrine like that proposed by this objection would be highly likely to do exactly that. Thus, the objection itself relies on a destabilizing and highly dangerous international norm. It may be the case that the AWS of the far future are to be treated as comparable to Marines, both tactically and in political terms, but anti-radiation missiles, loitering mortar rounds, and anti-missile gun batteries are clearly not to be deemed even remotely similar to flesh-and-blood combatants.³¹ As such, for the foreseeable future sending AWS will always be less escalatory than if we sent our own troops, and sending AWS may be able to tip the scales in favor of a defender, or at least increase their chances overall. Thus, AWS do provide a politically viable alternative to war.

4. Out of the Escalation Trap

The above talk of (de)escalation brings us to a further objection that has been raised against AWS, namely that their rapid and sometimes automatic responses to incoming

³⁰ Brown et al. (2022).

³¹ Many thanks to Maciej Zajac for posing this final objection to me.

information may make them highly escalatory, contrary to what I indicated above.³² The argument here is that AWS with specific mission parameters like, say, area denial, are liable to react to incursions into friendly space in a timeframe that does not admit any possibility for humans to prevent it, or even consider whether or not they should prevent it.³³ This could result in enemy units being killed when humans would not yet have intended that potentially escalatory response, and culminating in an escalation of the (incipient) conflict.³⁴ Relating to the overall argument of this article, the objection thus entails that AWS may in fact drag one into wars, rather than presenting an alternative to them.

Now, the first point to note is that even granting the premise of this objection – that the increased speed and possible automaticity of AWS may make them take action before humans could halt that – this does not actually present an objection to what we have argued for so far. This is because the objection moves the goalposts, as it were. The argument so far developed is that AWS present an alternative to war for would-be interveners or those hoping to help a state resist aggression. Thus, the departure point is a situation where a conflict is already underway, with a victim state being attacked by an aggressor or a population being subjected to grave human rights abuses (from either internal or external forces). However, the objection being raised concerns escalation between parties in a mounting, yet still non-kinetic conflict. Moreover, we have been discussing the provisioning of AWS by third parties as a way for those parties to make use of such systems as an alternative to war, whereas the objection focuses on degrading relations between potential pairs (or larger constellations) of belligerent actors. To that extent, the objection simply misses the mark.

For the sake of argument, however, let us look to a situation like that picked out by the objection, namely a case where one side of a mounting (but still non-kinetic) conflict deploys AWS. For concreteness, let us also return to our running example of the Russo-Ukrainian war. Supposing Ukraine had deployed advanced area denial AWS along its border prior to Russia's invasion, would the presence of such weapons have escalated the conflict in ways which otherwise would not have been the case? Clearly not, as Russia invaded even without such provocation. More weakly then, would such AWS have provided Russia with more legitimate grounds for military action, perhaps because those AWS engaged Russian units who had not yet crossed the border? Now, close proximity of adversarial military units can indeed lead to mistakes, sometimes even larger border skirmishes, either of which can obviously lead to broader escalation. However, there is no reason why AWS would necessarily carry out such actions. Ukraine could easily deploy area denial AWS along its border, and due to the AWS relying on exact GPS location, it could be made practically impossible for such systems to engage Russian units unless the

³² See, e.g., Altmann and Sauer (2017); Horowitz (2019); or Sauer (2020). Note also that some authors focus on the potential dire escalatory risks that would come with autonomous systems being involved in nuclear deterrence and response. See, e.g., Geist and Lohn (2018); Horowitz (2019); Johnson (2020); and Sauer (2020) pp. 249–252 (esp. the references included in footnote 71). However, as ably shown by Zajac (2022b) in this special issue, there are strong grounds to remove AWS from any aspect of nuclear deterrence and response. As such, we will focus on how AWS may impact on escalation in the realm of purely conventional warfare.

³³ Similar arguments could be made regarding counter-battery fire subsumed under autonomous systems, where incoming missiles or ordnance could trigger an automatic strike at the point of origin without a human even having the opportunity to review the decision.

³⁴ Wong et al. (2020) present this as a worry raised during a wargame conducted by the RAND corporation. In that wargame a missile strike meant to be only a demonstration of power led to automatic counterbattery fire from an AWS defense grid. Though it did not cause escalation in the game itself, such an action clearly has great escalatory potential.

latter had certainly crossed the border. As an added precaution, one could even deploy AWS a distance back from the border (say, twenty kilometers) to guarantee that no Russians who were merely lost or off course would be accidentally targeted. And even supposing there were Russians who were that lost, it is hard to see how utilizing force against enemy units twenty kilometers inside your own territory is any more escalatory than those enemy forces being that far into your territory in the first place.

This case can be made even stronger if we consider the fact that an AWS' targeting parameters may also be made public and transparent. Suppose Ukraine had anti-air, anti-armor, and anti-naval AWS deployed within its borders, and it publicly announced that all of these deployed systems would automatically engage any enemy who advanced more than twenty kilometers into Ukrainian territory, with no possibility for humans to override the AWS' engagement protocols. In that instance, the AWS would be, for all intents and purposes, no different from highly advanced and highly selective mines. Now, we may grant that it may be escalatory for Ukrainian soldiers to kill Russian soldiers who are only a short space within Ukrainian territory, but if Russian soldiers walk into clearly marked minefields, then the fault is only their own. Area denial AWS, if deployed with publicity and transparency, would be functionally no different than clearly marked minefields.

Given these points, AWS may arguably serve to be deescalatory, insofar as they allow a potential victim of aggression to set immediate and guaranteed consequences for a potential aggressor in the event of conflict. Moreover, the ability to publicly set targeting parameters allows for an open communication between adversary groupings regarding what will follow from certain actions. Putting certain autonomous systems into "full auto" mode and making it impossible for humans to override the AWS' actions can also present a way to show resolve,³⁵ an important aspect of deescalation in cases where one side will push until pushed back. And finally, AWS can serve to improve deescalation by providing a way out of the "escalation trap", i.e., situations where one side is willing to use another's escalation aversion as a way to exact political or territorial gains through the mere threat of escalation.³⁶ By setting uncompromising responses to aggression, potential aggressors will no longer be able to exploit the goodwill or humanitarian sentiments of rights-respecting regimes, and will instead have to calculate their actions against a backdrop of certain and possibly devastating robotic response.

5. Anti-Personnel AWS and the Weight of Dignity

All of the above arguments are essentially consequentialist in nature, focusing on how AWS may provide safer or more sure ways to respond to aggression and grave rights violations. However, some objections to AWS are not focused on the consequences of these systems, but rather argue that some essential feature of AWS makes them morally problematic, regardless of what may or may not follow from their adoption or use. One of the most common objections of this type is the contention that such systems would violate or undermine the dignity of those targeted by them, presenting a problem which cannot be mitigated by technical advances or other improvements on AWS' ability to be LOAC-compliant.³⁷ The remainder of this article will be devoted to showcasing a form of fully LOAC-compliant AWS which would run most afoul of this objection, and demonstrating that this potential (or actual) violation of dignity alone cannot suffice as a conclusive argument against AWS.

³⁵ Wong et al. (2020) p. xi, 52.

³⁶ See again Raine (2022), referenced above.

³⁷ See, e.g., Asaro (2012); Johnson and Axinn (2013); Sparrow (2016); Heyns (2016); Docherty et al. (2018); Rosert and Sauer (2021); or International Committee of the Red Cross (2021a). For a forceful rejection of the dignity-based objection to AWS, see Birnbacher (2016).

So far, our discussions have primarily focused on AWS which would target armor, aircraft, or naval vessels. There are a number of reasons for this, not least of which is that these systems are in many ways simpler than anti-personnel ones (and therefore more easily made to be compliant with the LOAC principle of distinction).³⁸ Moreover, with regards to our running example of the Russo-Ukrainian war, AWS of these exact types would have been incredibly helpful and provided many of the tactical capabilities the Ukrainian government most desperately wanted from Western countries providing aid. But let us now shift away from these sorts of systems and look instead to AWS which would explicitly target enemy combatants.

As a first point, it is worth noting that some objections to the use of anti-personnel AWS center on the idea that “robots of the near future will not be capable of making the sophisticated practical and moral distinctions required by the laws of armed conflict”³⁹ because in many cases determining whether or not someone is a legitimate target depends on the appraisal of subtle things like body language, tone of voice, or other nuanced aspects of human (non-verbal) communication.⁴⁰ While this objection will in some cases have bite, the rather obvious problem is that it only has bite *in some cases*. Moreover, the objection seems to be rooted in part in the fact “that the armed conflicts of the last decade and more have primarily been low-intensity counterinsurgency conflicts, and many commentators have unthinkingly taken the constraints that apply in such conflicts to be the norm for all forms of war”.⁴¹ However, if the conflict in question is not a counterinsurgency (with all of the complicating factors which attend that), but instead a regular war between regular units in regular uniforms, then this objection loses nearly all of its force. More concretely, if we are thinking about Ukrainian anti-personnel AWS targeting uniformed Russian combatants, then there are relatively straightforward targeting parameters which would make mistakes extremely unlikely while not greatly adversely affecting the effectiveness of such AWS. So let us begin by making exactly clear what we have in mind.

Since the main concern of the dignity-based objection is AWS which explicitly target humans (rather than those which target vehicles and weapons platforms, and only incidentally those humans who may be occupying those),⁴² that will be our locus of discussion: anti-personnel AWS. In order for such AWS to be at all permissible under the existing laws of war, they would have to be able to recognize a few key things as well. In particular, in addition to only targeting those with a clear combat dress of the enemy, they would have to be able to recognize when units are surrendering, or when units are *hors de combat*. These latter categories can be tricky, with many possibilities for mistakes and lethal errors. However, one can imagine rather straightforward ways to ensure compliance.

For example, if one only allowed anti-personnel AWS to target individuals who were i) wearing clear enemy combat dress, ii) were visibly armed, and iii) standing upright, then there is little possibility for mistakes to be made. By i), only enemies would be

³⁸ There are also independent moral and legal grounds for limiting AWS to only being able to target enemy vehicles or weapons platforms, or at least placing a moratorium on anti-personnel AWS. See, e.g., the arguments of Zajac (2022b) in this special issue.

³⁹ de Sio and Van den Hoven (2018) p. 2.

⁴⁰ Guarini and Bello (2012); Human Rights Watch (2016).

⁴¹ Baker (2022) p. 40.

⁴² Note that there are currently existing systems which count as AWS under the DoD/ICRC definition and which could kill humans, but which are not subject to any real critique by those seeking to ban AWS. For example, the CIWS turrets used by many navies can target incoming missiles *or high speed aircraft*. As such, it would seem that the core point of contention for those seeking to limit lethal AWS must be the targeting of humans specifically.

targeted. By ii), enemies can effectively signal an intent to surrender by simply dropping their arms, an action that takes the littlest of time. As a side effect, ii) would also make it so these AWS could only target enemies who present an actual threat, a level of care far in excess of what the LOAC necessarily demands of human combatants fighting, and greatly speaking in favor of such AWS from a humanitarian standpoint. And by iii), it would be virtually impossible for *hors de combat* combatants to be targeted, given that they would in virtually every case be unarmed, prone, or both.⁴³ At any rate, individuals who are wearing enemy dress, armed, and standing may obviously be targeted by human combatants at any time, and so there is no clearly immediate reason why an AWS would not likewise be able to target them under the LOAC. As a final precaution, we might further demand that such anti-personnel AWS fulfill only the lowest of tactical goals, and only in tightly defined areas of operation. They would therefore replace only “front-line grunts”, but never commanders or even non-commissioned officers, and they would always be linked to a specific location.⁴⁴ To coin a phrase, anti-personnel AWS would fulfill one purpose in modern regular war: replacing boots on the ground with bots on the ground.

We can imagine AWS like this, possibly even ones of the not-so-distant future, and though they would be LOAC-compliant they would still be, for all intents and purposes, eerily akin to *Terminators*; a robot would scour a battlefield looking for all active enemy combatants and, upon finding them, slaughter them without mercy. This is the science fiction horror that many opponents of AWS want us to imagine. But it is important to be clear about the fact that even these AWS can act within the letter of the law. It is also worth making clear that these AWS do not condemn all enemy combatants. Rather, the stringent targeting parameters sketched above mean that enemy combatants can make themselves safe by simply laying down arms or withdrawing outside the area of operations of the AWS. In the case of aggressor combatants, they need merely stop aggressing, either by dropping their weapons or returning to the country from which they came.⁴⁵ But let's imagine they don't do either of these things, and instead keep their weapons and continue their (unjust) fight.

In this case the AWS will target them, engage them, and most likely kill them.⁴⁶ According to the objection in question here, the AWS is argued to thereby infringe upon or violate the dignity of these combatants. But how exactly does the AWS cause this infringement or violation?

Recall that the example we have in mind is a Ukrainian anti-personnel AWS killing Russian combatants participating in the illegal and immoral invasion of Ukraine. Those combatants were almost certainly under at least some degree of duress when they were deployed to Ukraine, but they as individuals have protected rights to surrender, rights

⁴³ These parameters would expose the AWS to significant risks, especially risks of the parameters being abused by enemy forces (e.g., dropping arms long enough for the AWS to move on and then picking their weapons back up). These concerns would need to be addressed through new laws, general norms of war, or other guarantees. For the purposes of the arguments here, however, it is enough that these parameters make risks of mistakes extremely low.

⁴⁴ Again, similar concerns regarding AWS in mid- or higher-level command positions are ably discussed in Zajac (2022b).

⁴⁵ Thus, similar to the points made at the end of Section 4, anti-personnel AWS may also be able to serve a deescalatory function by inducing individual combatants to desert or surrender, weakening an aggressive country's ability to wage aggressive war.

⁴⁶ It is not clear that an AWS would have to lethally target, or that it even should (See, e.g., Wood (2020)). However, in order to present the most morally hazardous or objectionable case, we will focus on AWS which do by default make lethal targeting decisions.

which they may exercise at any point. Moreover, Ukraine could even follow the suggestions made in Section 4 above and make public the targeting parameters of their anti-personnel AWS in the hopes that this would induce Russian combatants to surrender. If the Russian soldiers still continue fighting, then this is a choice they are making, knowing full well what may follow. They know how to avoid being targeted, they know that carrying arms and continuing to be within Ukraine's territory makes their deaths likely, yet they do so anyway. In such a case, they are completely in control of the situation, vis-à-vis whether or not an AWS will target and kill them, and so their being killed by a robot is as much their choice as it is the robot's. So why do we focus solely on the "choice" the robot makes, and ignore the agency and humanity of the man who decides to enter the robot's area of operations, especially given that these sorts of AWS, when deployed with transparency, would be very similar to mines in marked minefields?

To make use of an even older analogy, did the sharpened stakes set out by English archers at Agincourt violate the dignity of those French knights who chose to charge and were impaled? No. The English set a hindrance to French movements, and one which would certainly kill some of the latter if they charged, and they charged anyway. They made a choice, a choice over which they had full control, and they paid a price. That is war, and it is hard to see how one side violates the dignity of another by modifying the choices available to the other. More generally, when one side has complete control over whether or not they will enter a highly hazardous area, and when they know what hazards lie in wait for them, it is difficult (to say the least) to see how these soldiers' dignity is violated because they knowingly walk into certain death, because they "charge the stakes", as it were. More strongly still, it seems a violation of these exact soldiers' dignity to *not give them their due responsibility* for their part in these things. Why are we to ignore the role they play, the choices they make, or their human agency? To do so is to treat them as less than full moral agents, which would seem one of the essential elements of violating one's dignity.

However, for the sake of argument, let us even grant the objection its thickest premise, namely that the use of AWS, any AWS, violates the dignity of those targeted. One may rejoin, "So what?" On those ethical theories that leave space for dignity, dignity is *a* moral concern. It is, however, not *the only* moral concern, nor one with overwhelming weight. In point of fact, on any sensible deontological moral theory, things like innocence, culpability, or liability to harm will arguably hold far more weight than mere dignity. After all, it would surely violate my dignity if I were drowned in sewage, but if I am trying to murder the innocent foreman of the sewage treatment plant and the only way to save him is to knock me into one of the murky pools (from which I cannot be saved), then this is certainly permissible.⁴⁷ So returning to the case of Ukrainian anti-personnel AWS killing Russian combatants, one must ask oneself, "So what?" We can assume it violates their dignity to be killed by a robot, but they are prosecuting an unjust war and committing a slew of war crimes along the way as well, including the killing of civilian men, women, and children. Moreover, it is arguably the case that they violate the dignity of every civilian who is killed as a result of their actions, and likely those Ukrainian combatants they kill as well (since they too are largely innocent, from a moral standpoint). So how or why is the dignity of unjust combatants so weighty as to present an argument against a weapon which could allow victims of aggression to repel that aggression? Maybe one can conjure some sort of argument to show how the dignity of an unjust combatant is indeed weightier than the value of an innocent child's life, but the onus of providing that argument rests with those hoping to base a ban on AWS in such dignity-related concerns. Until that argument is provided, however, we should view dignity as, at most, placing *a weight* in the scales against AWS. Whether or not that weight turns out to be decisive will

⁴⁷ Thanks to Maciej Zajac for suggesting this example to me.

vary case by case, and with regards to the heinousness of the aggression or rights abuses being perpetrated by those targeted by an AWS.

6. A Cautious Way Forward

In closing, it is worthwhile to stress that none of the arguments developed here should be construed as saying that all AWS are permissible, or that all AWS should be permissible, or that AWS should be used without care and thought about how they may alter the political, strategic, legal, or moral landscape of a given conflict. Quite on the contrary, the central point of this article is that there are and will continue to be a wide variety of weapon systems which count as autonomous (under the reigning definitions of the U.S. DoD or the ICRC), and that each individual system must be evaluated based on its own capabilities, limitations, and capacity to be compliant with the laws and moral norms of war. Some AWS will obviously present deep moral or legal problems, and will need to be limited, if not banned outright. However, with a clear appreciation of the variety of systems out there, it is also clear that some AWS will be permissible, both legally and morally. Moreover, these types of systems provide an important alternative to war or force short of war for would-be interveners or those otherwise hoping to help victims of aggression in repelling that unjust violence. Simpler AWS can also present a strong tool for deescalation, open a means for weaker states to place a check on the imperialistic ambitions of aggressive neighbors, or force the escalation decisions back onto those hoping to push rights-respecting regimes into an “escalation trap”.

With regards to concrete cases like the Russo-Ukrainian war, it is also clear that if we want to do the best we can to aid victims of aggression, but war is simply not an option (for political or moral reasons), then AWS may be able to provide a means to help with much lower risks of escalation. This is an important gain, both strategically and morally, and not to be eschewed simply because robots will be doing the killing for us. Moreover, even if there are grounds for taking a cautious approach to delegating kill-decisions to autonomous systems, we must ensure that this caution does not lead us to prioritize the humanity or dignity of unjust combatants over the preservation of innocent lives. War is hell, and that is not going to change anytime soon. The best we can hope for is to ensure that the innocent are spared and that the smallest amount of blood is shed.

Author Contributions: Nathan G. Wood is the sole author of this work.

Funding: This research received no external funding.

Acknowledgments: I must express my deepest thanks to Maciek Zajac, who discussed many of these ideas with me and helped push me to consider more incisive examples than those I otherwise would have. The work is far better for our weekly chats! And as always, I am deeply grateful to my wife Anna who always lets me air the first rough draft over dinner or while walking the dog together; your impact on this and all my other work cannot be overstated.

Conflicts of Interest: The author declares no conflict of interest.

References

- Altmann, J. and Sauer, F. (2017). Autonomous weapon systems and strategic stability. *Survival*, 59(5):117–142.
- Anderson, K. and Waxman, M. (2012). Law and ethics for robot soldiers. *Policy Review*, 176.
- Asaro, P. (2008). How just could a robot war be. In Briggie, A., Waelbers, K., and Brey, P. A., editors, *Current Issues in Computing and Philosophy*, pages 50–64. Ios Press, Amsterdam, Netherlands.
- Asaro, P. (2012). On banning autonomous weapon systems: Human rights, automation, and the dehumanization of lethal decision-making. *International Review of the Red Cross*, 94(886):687–709.
- Baker, D. (2022). *Should We Ban Killer Robots?* Polity.
- Barnett, M. N. (2002). *Eyewitness to a Genocide: The United Nations and Rwanda*. Cornell University Press, Ithaca, NY.
- Birnbacher, D. (2016). Are autonomous weapons systems a threat to human dignity? In Bhuta, N., Beck, S., Geiß, R., Liu, H.-Y., and Kreß, C., editors, *Autonomous Weapons Systems: Law, Ethics, Policy*, pages 105–121. Cambridge University Press, Cambridge, UK.
- Boothby, W. H. (2016). *Weapons and the Law of Armed Conflict*. Oxford University Press, Oxford, UK, 2 edition.

- Boulanin, V., Davison, N., Goussac, N., and Carlsson, M. P. (2020). Limits on autonomy in weapon systems: Identifying practical elements of human control. Technical report, International Committee of the Red Cross and Stockholm International Peace Research Institute.
- Brown, D., Horton, J., and Ahmedzade, T. (2022). Ukraine weapons: What military equipment is the world giving? *BBC News*. <https://www.bbc.com/news/world-europe-62002218>, Accessed: Sep. 25, 2022.
- CNBC Meets (2014). Interview with Bill Clinton on Rwanda. <https://www.youtube.com/watch?v=RpKV8psvnaQ>, Accessed: Sep. 22, 2022.
- Cooper, R. H. and Kohler, J. V., editors (2009). *Responsibility to Protect: The Global Moral Impact for the 21st Century*. Palgrave MacMillan, New York, NY.
- de Sio, F. S. and Van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, 5:1–14.
- Dinstein, Y. and Dahl, A. W. (2020). *Oslo Manual on Select Topics of the Law of Armed Conflict: Rules and Commentary*. Springer, Cham, Switzerland.
- Docherty, B. L., Duffield, D., Madding, A., and Shah, P. (2018). Heed the call: A moral and legal imperative to ban killer robots. Technical report, Human Rights Watch.
- Evans, G. (2009). The responsibility to protect: From an idea to an international norm. In Cooper, R. H. and Kohler, J. V., editors, *Responsibility to Protect: The Global Moral Compact for the 21st Century*, pages 15–30. Palgrave MacMillan, New York, NY.
- Geist, E. and Lohn, A. J. (2018). How might artificial intelligence affect the risk of nuclear war? Technical report, RAND Corporation, Santa Monica, CA.
- Guarini, M. and Bello, P. (2012). Robotic warfare: Some challenges in moving from noncivilian to civilian theaters. In Lin, P., Abney, K., and Bekey, G. A., editors, *Robot Ethics: The Ethics and Social Implications of Robotics*, pages 129–144. MIT Press, Cambridge, MA.
- Gubrud, M. (2014). Stopping killer robots. *Bulletin of the Atomic Scientists*, 70(1):32–42.
- Guichaoua, A. (2020). Counting the Rwandan victims of war and genocide: Concluding reflections. *Journal of Genocide Research*, 22(1):125–141.
- Heyns, C. (2013). Report of the special rapporteur on extrajudicial, summary, or arbitrary executions: UN doc. A/HRC/23/47. Technical report, UN General Assembly.
- Heyns, C. (2016). Autonomous weapons systems: Living a dignified life and dying a dignified death. In Bhuta, N., Beck, S., Geiß, R., Liu, H.-Y., and Kreß, C., editors, *Autonomous weapons systems: Law, ethics, policy*, pages 3–20. Cambridge University Press, Cambridge, UK.
- Horowitz, M. C. (2019). When speed kills: Lethal autonomous weapon systems, deterrence, and stability. *Journal of Strategic Studies*, 42(6):764–788.
- Human Rights Watch (2016). Killer robots and the concept of meaningful human control. <https://www.hrw.org/news/2016/04/11/killer-robots-and-concept-meaningful-human-control>, Accessed: June 17, 2022.
- International Committee of the Red Cross (2014). Autonomous weapons systems: Technical, military, legal and humanitarian aspects. Technical report, International Committee of the Red Cross.
- International Committee of the Red Cross (2021a). ICRC position and background paper on autonomous weapons systems. Technical report, International Committee of the Red Cross.
- International Committee of the Red Cross (2021b). ICRC position on autonomous weapons systems. Technical report, International Committee of the Red Cross.
- International Institute for Strategic Studies (2022). *The Military Balance 2022*. Routledge.
- Johnson, A. M. and Axinn, S. (2013). The morality of autonomous robots. *Journal of Military Ethics*, 12(2):129–141.
- Johnson, J. (2020). Delegating strategic decision-making to machines: Dr. Strangelove redux? *Journal of Strategic Studies*, 45(3):439–477.
- Kirby, J. and Guyer, J. (2022). Russia’s war in Ukraine, explained. *Vox*. <https://www.vox.com/2022/2/23/22948534/russia-ukraine-war-putin-explosions-invasion-explained>, Accessed: Sep. 7, 2022.
- Kofman, M. (2022). Why the Kremlin will seek regime change in Ukraine. *The Economist*. <https://www.economist.com/by-invitation/2022/02/23/michael-kofmanan-expert-on-russias-armed-forces-explains-why-the-kremlin-willseek-regime-change-in-ukraine>, Accessed: Sep. 22, 2022.
- Kofman, M., Migacheva, K., Nichiporuk, B., Radin, A., Tkacheva, O., and Oberholtzer, J. (2017). Lessons from Russia’s operations in Crimea and Eastern Ukraine. Technical report, RAND Corporation.
- Meierhenrich, J. (2020). How many victims were there in the Rwandan genocide? A statistical debate. *Journal of Genocide Research*, 22(1):72–82.
- PBS Frontline (1999). “The Triumph of Evil”, interview with Philip Gourevitch. <https://www.pbs.org/wgbh/pages/frontline/shows/evil/interviews/gourevitch.html>, Accessed: Sep. 22, 2022.
- Raine, J. (2022). Time for NATO to find a way out of the escalation trap in Ukraine. Technical report, International Institute for Strategic Studies. <https://www.iiss.org/blogs/analysis/2022/03/time-for-nato-tofind-a-way-out-of-the-escalation-trap-in-ukraine>, Accessed: Sep. 25, 2022.
- Roff, H. M. (2013). Killing in war: Responsibility, liability, and lethal autonomous robots. In Allhoff, F., Evans, N. G., and Henschke, A., editors, *Routledge Handbook of Ethics and War: Just War Theory in the 21st Century*, pages 352–364. Routledge, New York, NY.

- Roff, H. M. (forthcoming). An ontology of autonomy and autonomous weapons systems. In Finkelstein, C., MacIntosh, D., and Ohlin, J. D., editors, *The Ethics of Autonomous Weapons*. Oxford University Press, Oxford, UK.
- Rosert, E. and Sauer, F. (2021). How (not) to stop the killer robots: A comparative analysis of humanitarian disarmament campaign strategies. *Contemporary Security Policy*, 42(1):4–29.
- Sauer, F. (2020). Stepping back from the brink: Why multilateral regulation of autonomy in weapons systems is difficult, yet imperative and feasible. *International Review of the Red Cross*, 102(913):235–259.
- Sharkey, N. (2010). Saying “no!” to lethal autonomous targeting. *Journal of Military Ethics*, 9(4):369–383.
- Sharkey, N. (2012). Killing made easy: From joysticks to politics. In Lin, P., Abney, K., and Bekey, G. A., editors, *Robot Ethics: The Ethics and Social Implications of Robotics*, pages 111–128. MIT Press, Cambridge, MA.
- Sparrow, R. (2007). Killer robots. *Journal of Applied Philosophy*, 24(1):62–77.
- Sparrow, R. (2016). Robots and respect: Assessing the case against autonomous weapon systems. *Ethics & International Affairs*, 30(01):93–116.
- Strawser, B. J. (2010). Moral predators: The duty to employ uninhabited aerial vehicles. *Journal of Military Ethics*, 9(4):342–368.
- Toh, M., Ogura, J., Humayun, H., Yee, I., Cheung, E., Fossum, S., and Maruf, R. (2022). The list of global sanctions on Russian for the war in Ukraine. *CNN Business*. <https://edition.cnn.com/2022/02/25/business/list-globalsanctions-russia-ukraine-war-intl-hnk/index.html>, Accessed: Sep. 25, 2022.
- UK Ministry of Defence (2017). Joint doctrine publication 0-30.2: Unmanned aircraft systems. Technical report, United Kingdom Ministry of Defence.
- UN News (2022). General Assembly resolution demands end to Russian offensive in Ukraine. <https://news.un.org/en/story/2022/03/1113152>, Accessed: Sep. 25, 2022.
- United Nations Office of the High Commissioner for Human Rights (2022). Conflict-related civilian casualties in Ukraine. Technical report, UN Office of the High Commissioner for Human Rights.
- US Department of Defense (2017). DoD Directive 3000.09. Technical report, United States Department of Defense.
- Williams, A. P. (2015). Defining autonomy in systems: Challenges and solutions. In Williams, A. P. and Scharre, P. D., editors, *Autonomous Systems: Issues for Defense Policymakers*, pages 27–62. NATO Communications and Information Agency, The Hague, Netherlands.
- Wong, Y. H., Yurchak, J. M., Button, R. W., Frank, A., Laird, B., Osoba, O. A., Steeb, R., Harris, B. N., and Bae, S. J. (2020). Deterrence in the age of thinking machines. Technical report, RAND Corporation, Santa Monica, CA.
- Wood, N. G. (2020). The problem with killer robots. *Journal of Military Ethics*, 19(3):220–240.
- Wood, N. G. (forthcoming). Autonomous weapon systems: A clarification. *Journal of Military Ethics*.
- Wood, N. G. (unpublished manuscript). Rethinking proportionality. https://www.researchgate.net/publication/351658489_Rethinking_Proportionality.
- Zajac, M. (2022a). *Autonomous Weapon Systems from a Just War Theory Perspective*. PhD thesis, University of Warsaw.
- Zajac, M. (2022b). Beyond Deadlock; Low Hanging Fruit and Strict yet Achievable Options in AWS Regulation. *Journal of Ethics and Emerging Technologies*, 32(2):1–14.