

BURDEN OF PROOF IN THE AUTONOMOUS WEAPONS DEBATE

WHY BAN ADVOCATES HAVE NOT MET IT (YET)

Author: Maciek Zajac

Who has the burden of proof (BP) in the debate over the ethical permissibility of using autonomous weapon systems (AWS)?¹ This is undoubtedly an important question, especially given that this practical issue will eventually have to be resolved one way or the other. Unfortunately, few if any have so far attempted to answer it. This article aims to fill this gap in current research. I start with briefly showing why it is legitimate to treat the AWS debate as fit for differential BP assignment. I proceed to argue that BP should be assigned differently regarding various issues and classes of arguments within the AWS debate – falling on permissivists (opponents of a universal AWS ban) in regard to AWS capacity for compliance with the laws of war, and on prohibitionists (ban advocates) as far as other presumptive reasons for banning AWS are concerned. I also comment on how well the respective sides seem to shoulder BP assigned in this way.

Does the AWS debate need BP assignment? I believe it clearly does. The debate is deadlocked. On a diplomatic level, a decade of debate produced no new treaties or other universally binding legal documents, clearly frustrating prohibitionists², but also advocates of a basic regulatory program that has attracted support from both sides of the debate³. On the academic level, both sides are essentially stuck at their respective positions. Efforts aimed at producing compromise policy solutions, such as the just mentioned policy roadmap authored by Ronald Arkin, Stuart Russell, Paul Scharre and other prominent scholars⁴ failed to gather significant attention or following. No major participant in the debate has ever changed their mind about the basic issue. Representatives of both views tend to differ at the level of most basic intuitions⁵ or to question the other side's very motivations⁶. Clearly assigning BP, as well as enforcing other rules of honest and proper reasoning, could benefit all involved by providing incentive to deliver decisive results and to convince opponents, rather than just fellow travelers.

Abstract

The debate on the ethical permissibility of autonomous weapon systems (AWS) is deadlocked. It could therefore benefit from a differentiated assignment of the burden of proof. This is because the discussion is not purely philosophical in nature, but has a legal and security policy component and aims to avoid the most harmful outcomes of an otherwise unchecked development.

Opponents of a universal AWS ban must clearly demonstrate that AWS comply with the Law of Armed Conflict (LOAC). This requires extensive testing of specific models to verify compliance in realistic combat situations; working on appropriate testing regimes is also one of the most important future tasks for AWS ethics. In this context, however, it is easy to refute the claim put forward by prohibitionists that AWS violate LOAC a priori.

A state's right to ensure security and the protection of human rights through military capacity has a high ethical value. Other arguments against the use of AWS that are not based on LOAC conformity must therefore meet a substantial burden of proof, as they call on governments to abandon a military technology which is at least potentially transformative. None of those objections meet this standard. They include not only culturally or even individually divergent norms and ideas regarding the use of weapons, but also consequentialist (security policy) concerns on the one hand and so-called deontological arguments on the other. While the former at least point to the need for essential regulation, the latter are based on rather dubious premises. For the purposes of solutions beneficial to humanity, prohibitionist should formulate narrower and more specific points of criticism.

More importantly, the AWS debate is being waged on the clock – it is first and foremost a policy debate with immediate and massive real-life implications. A practical decision of some sort will inevitably follow even in the absence of a clear theoretical resolution – after all, AWS will either be developed and fielded worldwide or not. This is also a legal debate about what the international law is, and what it should be. As such it is precisely the type of a debate in which it is legitimate to assign BP to its respective sides.⁷ This is true even from a standpoint of the authors skeptical of the legitimacy of differential burden of proof assignment in philosophical debate in general, as the debate is clearly both about the law and about safety, that is, aims to elicit action that will avoid (the most) harmful outcomes⁸. Consequently, BP assignment would not only be helpful to get the AWS debate moving – the AWS debate is a paradigmatic example of a controversy that could and should benefit from its application.

Laws of Armed Conflict compliance

The Law of Armed Conflict (LOAC) represents a significant achievement of modern law – and of modern ethics. It applies to all weapon use – and so to all AWS use. LOAC is, in most of its aspects, a well-thought compromise between the requirements of military necessity and humanitarian concerns, and as such admits of no additional compromises⁹. Thus to argue that we need to relax or tighten its requirements because of some concern usually means that one is asking to double count that specific concern, as LOAC already does accommodate it. Accordingly, regarding LOAC compliance, it may seem that the burden of proof is staunchly on the permissivists – compliance with LOAC is not negotiable.

This definitely holds for specific models of AWS – specific instantiations of the general concept. Was any such weapon designed and manufactured, it should definitely be reviewed for compliance, which would include empirical “tests that realistically replicate the generic intended circumstances of use”,

including both the environment and intended mode of use¹⁰. Myself and many other permissivists readily agree that these tests should be appropriately rigorous and that no AWS should be deployed unless it can pass them¹¹.

I also suspect that many specific AWS designs, especially those featuring software manufactured using machine learning and analogous techniques, will not pass such

Participation in the design of appropriately rigorous and comprehensive testing regimes should become a focus, if not the focus, of AWS ethics

tests. Indeed, it is possible that, at least for a period of time, none of the proposed designs will. Such failures may generate pressures from both the weapons industry and interested militaries to lower the requirements. Insisting that predictability and reliability of AWS should reach desirable levels, and keeping the bar appropriately high, will be an important task for all ethicists dealing with AWS, if not the most important task indeed. Participation in the design of appropriately rigorous and comprehensive testing regimes should become a focus, if not the focus, of AWS ethics¹².

But this kind of skeptical empirical inquiry into potential flaws of specific AWS designs, or into the potential flaws of specific tactics involving these designs, is not what prohibitionists call for. Prohibitionists insist that such testing, or any development of AWS, should never occur, because according to (most) prohibitionists all AWS will of necessity be LOAC noncompliant, which can be known with certainty and *a priori*, only on the basis of the philosophical investigation of the very concept of AWS.

This, however, is simply and demonstrably not the case. AWS are an extremely broad weapon class, involving a potentially infinite number of designs capable of operating in all possible combat environments and scenarios. It is thus highly implausible to suggest

that no AWS whatsoever will ever be able to remain LOAC compliant in any possible environment or circumstances. If permissivists are to refute this claim, it is enough for them to provide just one example of an AWS that cannot be plausibly held to violate LOAC by operating in its intended environment, at least not *a priori*, based on purely philosophical considerations. Even if the burden of refuting the prohibitionist thesis stays on the permissivists, it is very easily met.

There are several examples of combat functions in specific operational environments that AWS could fulfill without risking LOAC violations, but let us focus on the one with arguably greatest military significance – achieving air superiority in high-end warfare. While NATO countries have been generally hesitant to develop AWS, the current US ‘loyal wingman’ efforts to involve large numbers of potentially autonomy-capable drones in air-to-air combat against a near-peer adversary are very substantial and seem to constitute the centerpiece of the US concept of air operations¹³. In essence, this class of (potential-

***If robust and reliable enough
AWS can be created, they are almost
certain to bestow on their users
a significant military advantage***

ly, optionally) autonomous weapon systems would be constituted by unmanned fighters capable of engaging aircraft within a zone of airspace closed to civilian traffic, and thus containing only military aircraft, whether friendly or enemy ones. Since aircrews have no method of rendering themselves *hors de combat* other than bailing out (there is no surrendering in a modern fighter or AWACS plane, especially in beyond visual range combat)¹⁴, this space by definition includes no non-combatants; no current occupant of a non-friendly aircraft entering this space can be wronged by being targeted. Thus as long as the loyal wingman AWS only attack air target within a designated combat zone, there is no conceivable reason to believe LOAC

violations would occur. Use outside of such zone, either due to a commander’s decision or a malfunction, could and most probably would be problematic – yet that is true of any weapon used inappropriately.

To summarize – permissivists have the burden of proof as far as AWS compliance with LOAC/*ius in bello* is concerned. Regarding the legality and ethicality of using any specific AWS design, this burden is substantial indeed; but in an *a priori* philosophical discussion it is very light and easily borne¹⁵. Indeed, some prominent prohibitionists have to their credit acknowledged this¹⁶ and attempted to move its focus towards other arguments, to which we now turn.

Applying the burden of proof to other prohibitionist arguments

Non-LOAC based arguments against AWS are diverse, yet they are all similarly positioned as far as the burden of proof assignment is concerned. As mentioned, LOAC-based arguments concern rules of conduct in war that already do incorporate considerations of military necessity – that is, legitimate concerns of national defense. Unless the rare conditions of the so-called “supreme emergency” obtain¹⁷, it thus makes no sense to say that one has to violate the laws of war because military necessity dictates it – this would lead to double counting this particular set of concerns. This, however, is not true regarding other kinds of arguments against AWS. Here considerations of military necessity can and should be weighed against the proposed downsides of AWS adoption.

AWS military value remains somewhat speculative, especially regarding specific types of AWS design, although it seems as certain as any prediction regarding the future of technology can be. If robust and reliable enough AWS can be created – a matter to be determined empirically – they are almost certain to bestow on their users a significant military advantage. Simply put, AWS offer a way to transcend all human limitations in tactical combat. Cheaper, easier to deploy,

more accurate, faster, with better readiness levels ultimately expandable yet promising perfect skill retention across the force – AWS would synergistically combine all these features, overpowering human combatants or drone operators. While these characteristics need not be delivered by every single AWS design, they seem to be a logical outcome of any competent, well-resourced, long-term research and development effort, just as surpassing horse-drawn carriages in speed was a logical outcome of efforts to create automobiles. For the purposes of this article, let us assume that AWS would indeed bestow a significant military advantage, at the very least comparable to the one bestowed by the development of drones. I am comfortable with making this assumption for several reasons. Most importantly, unless we make it, the entire AWS debate loses its salience.¹⁸

Offering a military advantage, even a revolutionary one, is not an ethical property. However, a choice to reject such an advantage is ethically loaded, given that both military and civilian authorities have a duty to effectively defend their citizens against plausible security threats. Rejecting a revolutionary military technology *without a substantial reason* violates that duty. That AWS' putative inability to comply with LOAC would furnish such a reason, is, at least *prima facie*, uncontroversial – LOAC is a product of well-established moral consensus. Alternative reasons to reject AWS should be comparably uncontroversial – grounded in commonly acknowledged values and based in sound assumptions. Prohibitionists are required to furnish such reasons. To realize why no lesser argument will do, one has to understand how ethically fundamental the government's duty is to protect its citizens.

Within Western nations with few obvious security concerns it is somewhat fashionable to view the nation's security apparatus as an organization akin to a tobacco company – a nefarious institution exploiting human vice to stay in business. Consequently, decreases in a nation's military capacity may be regarded positively, and any increase with suspicion. Such an analysis is nevertheless a profound-

ly unserious attempt to describe the political, social and moral role of a security apparatus in a liberal democracy. Such a state, and its judicial and executive apparatus, is the primary and exclusive vehicle for realizing its inhabitants' basic human rights. These paradigmatically essential rights are recognized by the broadest ethical consensus to have ever existed globally as absolutely essential to human flourishing of any kind¹⁹. However, human rights cannot be realized – protected and delivered – but through a set of institutions that have to jointly exercise a number of capacities that we know as attributes of state sovereignty, such as the monopoly on violence or right to impose taxation²⁰. Any

Within Western nations with few obvious security concerns it is somewhat fashionable to view the nation's security apparatus as an organization akin to a tobacco company

violation of sovereignty directly and severely impacts on the state's capacity to safeguard and promote basic human rights, and consequently violates these very rights. This imbues the capacity to enforce and retain sovereignty over the state's territory – military capacity – with very substantial ethical value, provided it is exercised by a government that is not derelict in its duty to respect and promote human rights.

This capacity is all the more essential to governments that are facing active threats to their sovereignty. Ukraine, Taiwan and South Korea are three paradigmatic examples of a government – and society – placed in this situation. It is abundantly clear that only these government's military capacity, as well as the military capacity of their allies, prevents these governments from being destroyed, and their citizens from being subjected to severe and universal violations of their human rights. To tell the Ukrainians, the Taiwanese or the South Koreans not to develop a given military capacity is to ask them to place themselves at substantial risk – and so the burden of proof is on the person making the demand.

Nor are the circumstances of Ukraine, Taiwan or South Korea exceptional. Many other countries do deter plausible military threats only, or mostly, through their military capacity. That this is not so immediately obvious is caused by this capacity being extremely robust, not by the absence of such threats. Still other countries owe their peace, in part or fully, to the protective efforts of their allies. When being asked not to pursue a potentially transformative military technology, all of these countries are being asked to take a risk – and this sacrifice needs to be offset by a powerful enough reason.

It is important to note that to assign – or assume – this burden is not to prejudice the discussion in any way, just as permissivists assuming their share of the burden vis-à-vis LOAC compliance does not unfairly handicap them. What it does, however, is require the arguments put forward to be well-crafted, specific enough, valid and based on universally accepted premises, including axiological ones. A right to pursue security through

A right to pursue security through military capacity is a universally recognized right. Reasons invoked against it need to share this quality

military capacity is a universally recognized right. Reasons invoked against it need to share this quality.

Another feature of serious attempts to argue against moral feasibility of AWS use should be employment of criteria that can be applied to all other weapon classes²¹. LOAC standards constitute a paradigmatic example of these. A weapon class's impact on strategic stability can be another basis for ethical evaluation – “would this weapon cause unmanageable disruption?” is a question frequently and sensibly asked about various systems. On the other hand, “can this weapon be wielded honorably?” is currently being asked of AWS²², but not of other weapon systems; asking this question thus leads to an application of a double standard, as well as oth-

er predictable problems. Since this criterion is not a product of universal consensus and continued practice, it is very likely to import culturally dependent norms, or even individual idiosyncrasy, into the debate.

To summarize: non-LOAC-based anti-AWS arguments have to meet a quite substantial burden of proof, since they urge governments facing serious military threats to abandon a military technology with a potentially transformative impact. Such arguments need to invoke values and concerns that are as universal and substantial as the concerns of national defense, and that can simultaneously be invoked in relation to other weapon classes. Is this the case for other prohibitionist arguments?

Heavy burden, wobbly foundations

Arguments against AWS can be broadly classified into three groups²³, one of which we already discussed. The second group focuses on long-term problems that could be engendered by widespread AWS proliferation, such as the worry that this will increase the global incidence of armed conflict, or enable either terrorists and/or authoritarian states.²⁴ The third group of arguments, in contrast, can be called “non-consequentialist” (Sharkey terms them “deontological arguments”). These postulate AWS use would be detrimental in ways that are not reflected by worries about their effects on the security and well-being of persons as these are ordinarily understood. In this final section I will briefly discuss how well these two classes have been able to shoulder their due burden of proof.

The four most substantial strategic concerns are the worry about AWS adoption's impact on the frequency of armed conflict²⁵; their effect on global strategic stability, including the nuclear balance²⁶; the adverse effects of their proliferation to non-state actors, including terrorist groups²⁷; and the adverse effects of their use by authoritarian and totalitarian regimes²⁸. While varied, these arguments all have as their object the same

values as LOAC norms and the pro-AWS argument from a right to national defense – they aim to prevent tangible harm to persons in the form of physical violence or oppression grounded in a threat of violence. Whether these concerns are valid or not, they are definitely centered on salient objectives – there is no telling someone distressed about nuclear stability or increased incidence of terrorism that these matters are not important enough to be raised in the AWS debate. Thus these strategic concerns, both individually and as a class, clearly meet the meta-standards of specificity and saliency discussed in the previous section.

What they struggle to meet is the accompanying standard of validity. The nuclear stability argument fares best, as the scenarios of AWS being intentionally or inadvertently used to target nuclear deterrent forces are both plausible and deeply concerning. What is doubtful is the necessity of enforcing a universal AWS ban to address this concern, given that only certain types of AWS would pose a threat to nuclear deterrents.

This problem gets much worse for the twin arguments concentrating on proliferation to terrorist and authoritarian regimes. Yes, such outcomes would be disastrous – yet it seems that introducing a universal ban would be neither necessary nor sufficient to prevent them. Both terrorist attacks and authoritarian oppression are, after all, already illegal under current international law. Passing another treaty will do little to change the behavior of actors who by definition already reject legal and ethical norms, and are successful in evading efforts to enforce these norms. There is also no causal connection between legitimate governments building – or refusing to build – LOAC-compliant high-end AWS useful for addressing their security concerns, and authoritarians or terrorists producing low-tech killer robots exclusively useful for attacking defenseless civilians.

As for concerns about the increased incidence of war, these seem to assume casualty sensitivity as the key restraining factor; in reality, several other factors restrain actors with high casualty sensitivity, while low casualty

sensitivity actors would remain unaffected. It may also be argued that it is not the incidence of armed conflict, but the incidence of interstate aggression and intra-state oppression that should be the focus of concern, with AWS use enabling currently fledgling responses to these grim phenomena.

Discussing either of these strategic concerns in detail, let alone all of them, is beyond the scope of this paper – all I can do here is indicate some serious problems with referring to these generally legitimate wor-

Passing another treaty will do little to change the behavior of actors who by definition already reject legal and ethical norms, and are successful in evading efforts to enforce these norms

ries to argue for a specific policy solution, i.e., the AWS ban. While focusing on this class of arguments seems the most promising prohibitionist approach available, the burden of proof has not been met yet by these frequently underdeveloped arguments.

That said, consequentialist arguments against AWS appear robust and well-crafted when contrasted with their non-consequentialist counterparts. These include criticisms of putative arbitrariness of AWS targeting²⁹, calls to preserve a semblance of fairness in combat or warfare³⁰, and assertions that AWS use would violate the dignity of targeted enemy combatants. The first argument is the weakest. It either presupposes AWS inability to comply with LOAC, in which case it can be reduced to a LOAC-based argument, or it does not; if the latter, we are left with a puzzling claim that a LOAC compliant attack is nonetheless arbitrary, even though an attack targeting a combatant because he is a combatant is discriminate *ex hypothesi*. One could go on and express anguish at the thought of fellow humans being targeted just because they are part of an organized military endeavor to kill and subjugate others – but this would be an argument against war as such, by no means exclusive to AWS warfare.

The argument from fairness is, ironically, similarly unfair – and it is also similarly misguided. If one is fighting in justifiable defense of oneself or others, offering the enemy a fair fight is not only not obligatory, it is morally wrong. And if one does not have a good justification to fight, the act is wrong regardless of whether the fight is fair or not. An assault does not become justified because the attacker and the victim are in fact evenly matched; a victim of aggression is not required to allow the perpetrator any chance of victory.

As for the various dignity-based arguments, even though frequently voiced³¹, these fail to precisely articulate in what way LOAC-compliant AWS would violate the dignity of targeted enemy combatants³². To treat enemy combatants in accordance with LOAC

enemy combatants that would be violated by targeting them with AWS is specified, the assertion that their human dignity is violated remains unsupported. It is also hard to shake the impression that the intuitions about combatants' dignity being violated are engendered by considering scenarios in which LOAC is being violated, or by imagining relationships and attitudes towards each other that combatants do not have and could not be required to have³⁴.

Conclusion

The case of non-consequentialist arguments demonstrates why it is important to assign the burden of proof in the AWS debate. As in every debate, it is possible to advance a number of insufficiently grounded and increasingly idiosyncratic arguments. Rejecting those *a priori* would be wrong; one cannot know in advance whether an argument that appears idiosyncratic at first glance may in fact prove successful. Yet taking the very existence of miscellaneous appeals to non-universal values or arguments resting on shaky premises as evidence of a powerful ethical case against AWS is equally wrong.

Permissivists have to shoulder the considerable burden of proving that any specific AWS design can be used in LOAC-compliant manner, and of creating testing & verification methodologies capable of conclusively deciding the issue. This they have yet to achieve, especially regarding the more ethically difficult combat environments. They also have to make a plausible case that LOAC compliance is theoretically possible in order to justify such research. This I think they have already done.

Outside of LOAC compliance issues, permissivists can point to plausible benefits of AWS development, expressing them in the currency of universally accepted values and rights. Prohibitionists have to provide equally strong reasons against AWS use. While their consequentialist arguments have so far failed to do so, they do invoke concerns of considerable gravity that require regulatory responses. In contrast, non-consequentialist

Non-consequentialist arguments against AWS are either flawed, reducible to other types of arguments or fail to validly connect AWS use to specific, tangible harm

is already to respect their intrinsic worth as human beings (*Menschenwürde*) as much as possible given that they are actively participating in efforts to kill their opponents. To postulate that one should go beyond respecting LOAC, for example by showing mercy³³, is to postulate one should benefit enemy combatants at the expense of putting other people, including non-combatants, in harm's way. This is morally unacceptable. On the other hand, if no specific duty towards

The Author



Maciek Zajac, currently a postdoc on a Polish National Science Center "Sonatina" grant, received his PhD in philosophy from the University of Warsaw in 2022. His dissertation was entitled "Autonomous Weapon Systems in Just War Theory Perspective". He has published on the ethical issues surrounding AWS, as well as on other topics within military ethics in "Ethics and Information Technology", "Journal of Military Ethics", "Diametros" and other journals.

arguments against AWS are either flawed, reducible to other types of arguments or fail to validly connect AWS use to specific, tangible harm.

It is hard to escape the conclusion that more specific prohibitionist critiques of narrower scope pack more power. Prohibitionists would do best to work on those; permissivists would do best to listen to these carefully. Moving the debate towards offering feasible solutions beneficial to all humanity is, after all, a goal we all share.

1 ICRC has defined AWS as robotic weapon systems that “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 ‘target profile.’” I will follow this definition. International Committee of the Red Cross (2021): ICRC position on autonomous weapon systems and Background Paper, May 17th. <https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systems>. (All internet sources accessed April 22, 2024.)

2 Amoroso, Daniele (2020): Autonomous Weapons Systems and International Law: A Study on Human-Machine Interactions in Ethically and Legally Sensitive Domains. Naples, pp. 250-256; Jones, Isabelle (2021): Historic opportunity to regulate killer robots fails as a handful of states block the majority. Campaign to Stop Killer Robots press release, December 17th. <https://www.stopkillerrobots.org/news/historic-opportunity-to-regulate-killer-robots-fails-as-a-handful-of-states-block-the-majority/>; Moyes, Richard (2019): Critical Commentary on the ‘Guiding Principles’. Article 36, November. <https://article36.org/wp-content/uploads/2019/11/Commentary-on-the-guiding-principles.pdf>.

3 Arkin, Ronald et al. (2019): Autonomous Weapon Systems, a Roadmapping Exercise. Georgia Tech Robot Laboratory Online Publications, September 9th. <https://www.cc.gatech.edu/ai/robot-lab/online-publications/AWS.pdf>.

4 See endnote 3.

5 Baker, Deane (2022): Should We Ban Killer Robots? Cambridge, Medford, MA, pp. 96-99.

6 Sometimes publicly (cf. Sparrow, Robert (2023): A military-philosophical complex. In: *Metascience* 32, pp. 421-424. <https://doi.org/10.1007/s11016-023-00902-4>), more often off the record.

7 “It is the link to specific practical outcomes that is what the burden of proof is all about”: Hahn, Ulrike, and Mike Oaksford (2007): “The burden of proof and its role in argumentation.” In: *Argumentation* 21, pp. 39-61, p. 42. “We must frequently make decisions and act on the basis not of conclusive evidence but of what is reasonable to presume true”: Räikkä, Juha (2005): Global Justice and the Logic of the Burden of Proof. In: *Metaphilosophy* 36.1-2, pp. 228-239, p. 228.

8 Dare, Tim, and Justine Kingsbury (2008): Putting the burden of proof in its place: When are differential allocations legitimate? In: *The Southern Journal of Philosophy* 46.4, pp. 503-518, pp. 503-509.

9 Dinstein, Yoram (2016): *The Conduct of Hostilities under the Law of International Armed Conflict*. Third Edition. Cambridge, pp. 8-10; Melzer, Nils (2016): *International Humanitarian Law: a Comprehensive Introduction*. Geneva: International Committee of the Red Cross, pp. 17-18.

10 Boothby, William H. (2014): *Conflict law: the influence of new weapons technology, human rights and emerging actors*. Berlin, Heidelberg, pp. 168-176.

11 Zając, Maciej (2022): *Autonomous Weapon Systems from Just War Theory perspective*. PhD Thesis, University of Warsaw, pp. 158-161; Anderson, Kenneth, and Matthew C. Waxman (2017): *Debating Autonomous Weapon Systems, their Ethics, and their Regulation under International Law*. In: Brownsword, Roger, Eloise Scotford, and Karen Yeung (eds.): *The Oxford Handbook of Law, Regulation and Technology*. Oxford, pp. 1097-1117; Burri, Susanne (2018): What is the moral problem with killer robots. In: Strawser, Bradley, Ryan Jenkins and Michael Robillard (eds.): *Who Should Die? The Ethics of Killing in War*. New York, pp. 163-185, p. 178; Lin, Patrick, George Bekey, and Keith Abney (2008): *Autonomous military robotics: Risk, ethics, and design*. California Polytechnic State University San Luis Obispo, p. 71; Lucas Jr, George R. (2011): *Industrial challenges of military robotics*. In: *Journal of Military Ethics* 10.4, pp. 274-295; Wood, Nathan Gabriel (2020): *The problem with killer robots*. In: *Journal of Military Ethics* 19.3, pp. 220-240, p. 221.

12 For preliminary attempts to lay foundations for such work see Haugh, Brian A., David A. Sparrow, and David M. Tate (2018): *Status of Test, Evaluation, Verification, and Validation (TEV&V) of Autonomous Systems*. Institute for Defense Analyses. <https://apps.dtic.mil/sti/trecms/pdf/AD1118676.pdf>; Meier, Michael W. (2016): *Lethal autonomous weapons systems (laws): conducting a comprehensive weapons review*. In: *Temple International and Comparative Law Journal* 30, pp. 119-132; Verbruggen, Maaike (2022): No, not that verification: Challenges posed by testing, evaluation, validation and verification of artificial intelligence in weapon systems. In: Reinhold, Thomas and Niklas Schörnig (eds.): *Armament, Arms Control and Artificial Intelligence: The Janus-faced Nature of Machine Learning in the Military Realm*. Cham, pp. 175-191.

13 Trevithick, Joseph (2024): *Everything New We Just Learned About The Collaborative Combat Aircraft Program*. In: *The War Zone*, February 23rd. <https://www.twz.com/air/collaborative-combat-aircraft-poised-to-reshape-the-air-force>.

14 Surrender of aircraft by radio (Dinstein, Yoram (2016), see endnote 9, pp. 133-34) or by other methods (De Preux, Jean (1987): Article 41 – Safeguard of the enemy hors de combat. In: Pilloud, Claude et al. (eds.): *Commentary on the additional protocols of 8 June 1977 to the Geneva Conventions of 12 August 1949*. Geneva, pp. 479-492, p. 487) may be practicable of interdicted passenger or transport aircraft, but not in high-end BVR jet combat.

15 The reason why so many prohibitionists fail to see this is undue focus on ethically complex counterinsurgency scenarios to the exclusion of various other environments and combat tasks (Brenneke, Matthias

- (2020): Lethal Autonomous Weapon Systems and Their Compatibility with International Humanitarian Law: A Primer on the Debate. In: *Yearbook of International Humanitarian Law*, Volume 21 (2018). The Hague, pp. 59-98, p. 68.; Foy, James (2014): Autonomous weapons systems: Taking the human out of international humanitarian law. In: *Dalhousie Journal of Legal Studies* 23, pp. 47-70, p. 55; Geiss, Robin, and Henning Lahmann (2017): Autonomous weapons systems: a paradigm shift for the law of armed conflict? In: Ohlin, Jens David (ed.): *Research Handbook on Remote Warfare*. Cheltenham, Northampton, MA, pp. 371-404, p. 395.
- 16 Sparrow, Robert (2016): Robots and respect: Assessing the case against autonomous weapon systems. In: *Ethics & International Affairs* 30.1, pp. 93-116, p. 105; Rosert, Elvira, and Frank Sauer (2021): How (not) to stop the killer robots: A comparative analysis of humanitarian disarmament campaign strategies. In: *Contemporary Security Policy* 42.1, pp. 4-29, p. 14.
- 17 Orend, Brian (2013); *The Morality of War*. Second Edition Expanded and Updated. Peterborough, pp. 153-171; Walzer, Michael (2006): *Just and Unjust Wars: A Moral Argument with Historical Illustrations*. Fourth Edition. New York, pp. 251-268.
- 18 I believe the underlying logic of this assumption – and so do others (Collins, Liam and Harrison “Brandon” Morgan (2020): *Affordable, Abundant and Autonomous: The Future of Ground Warfare*. War on the Rocks, April 21st. <https://warontherocks.com/2020/04/affordable-abundant-and-autonomous-the-future-of-ground-warfare/>; Etzioni, Amitai, and Oren Etzioni (2017): Pros and Cons of Autonomous Weapons Systems. In: *Military Review* May-June, pp. 72-81, pp. 72-73; Hurst, Jules (2017): Robotic swarms in offensive maneuver. In: *Joint Force Quarterly* 87.4, pp. 105-111; O’Neill, Paul, Sam Cranny-Evans and Sarah Ashbridge (2024): *Assessing Autonomous Weapons as a Proliferation Risk*. Royal United Services Institute, February, p. 7. <https://static.rusi.org/future-laws-occasional-paper-feb-2024.pdf>; Watts, Sean (2016): Autonomous weapons: Regulation tolerant or regulation resistant. In: *Temple International and Comparative Law Journal* 30, pp. 177-186, pp. 178-80). In addition, I am unaware of it ever being challenged by a prominent prohibitionist publication.
- 19 Orend, Brian (2002): *Human Rights: Concept and Context*. Peterborough, pp. 33-34. In a famous phrase, humans want those rights no matter what else they want (Nickel, James (2008): *Making Sense of Human Rights*. Second Edition. Malden, pp. 10-11.)
- 20 Buchanan, Allen E. (2004): *Justice, legitimacy, and self-determination: Moral foundations for international law*. Oxford.
- 21 Baker, Deane (2022), see endnote 4, pp. 45-46; Heller, Kevin (2023): The Concept of ‘The Human’ in the Critique of Autonomous Weapons. In: *Harvard National Security Journal* 15, pp. 1-76, pp. 9-17.
- 22 Johnson, Aaron M., and Sidney Axinn (2013): The morality of autonomous robots. In: *Journal of Military Ethics* 12.2, pp. 129-141, pp. 135 f.
- 23 Sauer, Frank (2020): Stepping back from the brink: Why multilateral regulation of autonomy in weapons systems is difficult, yet imperative and feasible. In: *International Review of the Red Cross* 102.913, pp. 235-259; Sharkey, Amanda (2018): Autonomous weapons systems, killer robots and human dignity. In: *Ethics and Information Technology* 21.2, pp. 75-87, p. 78.
- 24 Frank Sauer calls these “strategic implications”, while Amanda Sharkey uses “consequentialist reasons”.
- 25 Sauer, Frank (2020), see endnote 23, pp. 249-251.
- 26 Sauer, Frank (2020), see endnote 23, pp. 251-252.
- 27 Kwik, Jonathan (2022): Mitigating the Risk of Autonomous Weapon Misuse by Insurgent Groups. In: *Laws* 12.1, p. 5. <https://doi.org/10.3390/laws12010005>.
- 28 Asaro, Peter (2012): On banning autonomous weapon systems: human rights, automation, and the dehumanization of lethal decision-making. In: *International Review of the Red Cross* 94.886, pp. 687-709. The two latter concerns, although being some of the most prominent and earliest ones, have been subject to alarmingly little focused research.
- 29 Asaro, Peter (2012), see endnote 28, pp. 697-700; Heyns, Christof (2017): Autonomous weapons in armed conflict and the right to a dignified life: an African perspective. In: *South African Journal on Human Rights* 33.1, pp. 46-71, pp. 57-60.
- 30 Johnson, Aaron M., and Sidney Axinn (2013), see endnote 22; Killmister, Suzy (2018): Remote weaponry: The ethical implications. In: *Journal of Applied Philosophy* 25.2, pp. 121-133.
- 31 Heyns, Christof (2016): Autonomous weapons systems: Living a dignified life and dying a dignified death. In: Bhuta, Nehal et al. (eds.): *Autonomous Weapons Systems: Law, Ethics, Policy*. Cambridge, pp. 3-20; Docherty, Bonnie Lynn, et al. (2018): Heed the Call: A Moral and Legal Imperative to Ban Killer Robots. *Human Rights Watch*. https://www.hrw.org/sites/default/files/report_pdf/arms0818_web.pdf; Rosert, Elvira, and Frank Sauer (2019), see endnote 16; Sparrow, Robert (2016), see endnote 16, pp. 110-112.
- 32 Birnbacher, Dieter (2016): Are Autonomous Weapons Systems a Threat to Human Dignity? In: Bhuta, Nehal et al. (eds.), see endnote 31, pp. 105-121, pp. 105-116.
- 33 Zajac, Maciej (2020): No Right To Mercy – Making Sense of Arguments From Dignity in the Lethal Autonomous Weapons Debate. In: *Etyka* 59 (1), pp. 134-155
- 34 Heller, Kevin (2023), see endnote 21, pp. 14-17.