




Autonomous weapon systems and *jus ad bellum*

Alexander Blanchard¹ · Mariarosaria Taddeo^{1,2} 

Received: 30 November 2021 / Accepted: 24 February 2022 / Published online: 19 March 2022
© The Author(s) 2022

Abstract

In this article, we focus on the scholarly and policy debate on autonomous weapon systems (AWS) and particularly on the objections to the use of these weapons which rest on *jus ad bellum* principles of proportionality and last resort. Both objections rest on the idea that AWS may increase the incidence of war by reducing the costs for going to war (proportionality) or by providing a propagandistic value (last resort). We argue that whilst these objections offer pressing concerns in their own right, they suffer from important limitations: they overlook the difficulties of calculating *ad bellum* proportionality; confuse the concept of proportionality of effects with the precision of weapon systems; disregard the ever-changing nature of war and of its ethical implications; mistake the moral obligation imposed by the principle of last resort with the impact that AWS may have on political decision to resort to war. Our analysis does not entail that AWS are acceptable or justifiable, but it shows that *ad bellum* principles are not the best set of ethical principles for tackling the ethical problems raised by AWS; and that developing adequate understanding of the transformations that the use of AWS poses to the nature of war itself is a necessary, preliminary requirement to any ethical analysis of the use of these weapons.

Keywords Artificial intelligence · Autonomous weapon systems · Jus ad bellum · Last resort · Proportionality · Just war theory

1 Introduction

Whilst the use of autonomous weapon systems (AWS) that are able to “identify, select and attack the target” without human intervention¹ looks increasingly likely in modern warfare, the ethical and legal challenges that they pose are far from being addressed (Taddeo et al. 2021). That is not to say that

debate on these implications have not been extensive. Since the publication of an executive order by the US Department of Defence on AWS (2012), there has been active and sustained consideration on the ethics of AWS by multiple actors, including policy-makers, roboticists, academics, civil society organizations, and state actors, particularly at the United Nation’s Convention on Certain Conventional Weapons Group of Governmental Experts related to emerging technologies in the area of lethal autonomous weapon systems (LAWS). The use of AWS poses a number of pressing ethical questions, and the debate has taken place along a number of fronts accordingly, from the locus of responsibility if or when the use of AWS goes wrong (Sparrow 2007; Schulzke 2013; Neha Jain 2016; Himmelreich 2019; Taddeo and Blanchard Forthcoming), to the loss of dignity implied by the decision to delegate the decision to kill to artificial agents (Birnbacher 2016; Heyns 2017; Sharkey 2019), to the nature of meaningful human control over AWS (Amoroso and Tamburrini 2020).

Above all, one area of particular concern is whether AWS can be used in a way which complies with the principles of international humanitarian law (IHL)—proportionality, discrimination, and necessity—which are comprised of a number of existing conventions on the conduct of war including the Geneva and Hague conventions (Sharkey 2010; Grut 2013;

¹ Taddeo and Blanchard (2021, 18) provide a full definition of AWS: “...an artificial agent which, at the very minimum, is able to change its own internal states to achieve a given goal, or set of goals, within its dynamic operating environment and without the direct intervention of another agent and may be endowed with some abilities for changing its own transition rules without the intervention of another agent, and which is deployed with the purpose of exerting kinetic force against a physical entity (whether an object or a human being) and to this end is able to identify, select and attack the target without the intervention of another agent is an AWS. Once deployed, AWS can be operated with or without some forms of human control (in, on or out the loop).”

✉ Mariarosaria Taddeo
mariarosaria.taddeo@oii.ox.ac.uk

¹ Alan Turing Institute, London, UK

² Oxford Internet Institute, University of Oxford, Oxford, UK

Sassoli 2014; Foy 2014; Chengeta 2016; Convention on Certain Conventional Weapons 2019). This is not just a question of technical capability, about whether in the future more refined models of AWS will be capable of contextualized judgement in a way that, say, allows them to distinguish combatants from non-combatants. The challenge that AWS have adhering to IHL is more fundamental than that, for these systems have novel capabilities (e.g. adapting capabilities) and enable new modes of operations (e.g. autonomy of the weapon systems) that can breach the foundational tenants underpinning IHL.

Given the profundity of this challenge scholars have begun to return to Just War theory, the philosophical tradition in which much of IHL is rooted to reconceptualize the novel tensions posed by AWS to the laws of war (Keith Abney 2013; Demy 2020; Roach and Eckert 2020). In this article, we challenge two objections that have emerged out of the application of Just War principles to AWS: that of *ad bellum* proportionality and last resort.

Just War theory is comprised of two sets of principles for governing the conduct of war: *jus ad bellum*—the principles that will concern us in this article—and *jus in bello*. The latter establishes the principles for right conduct in war, principles such as proportionality, distinction, and necessity (Coverdale 2004, 260–75). *Jus ad bellum* comprises the reasons for going to war, it establishes four conditions which must be met in order for a war to be considered justified as a whole. The first is that war must be conducted for a just cause. The second condition is that war must be declared by a legitimate authority. The third is that the costs of the war—being understood broadly—must be proportionate to the goals sought. The fourth and final condition for a war to be just is that it must be undertaken as a last resort—i.e. once all other forms of diplomacy and punitive measures have been exhausted (Coverdale 2004, 229–60).²

Since Just War theory holds a strict separation between *jus ad bellum* and *jus in bello* considerations, and since the use of weapons is thought to fall, almost exclusively, under *in bello* considerations, we would expect there to be little scope for applying *jus ad bellum* considerations to AWS (Walzer 1977, 21, 41–44).³ For instance, that war must be initiated by a lawful authority is a demand which, according to Just War

theory, is unrelated to the weapons employed in the course of the war. ‘Just cause’, too, is a moral and political question established on the basis of the given contingencies leading up to war (Christopher Finlay 2019, 33–71). It is, then, a sign of the potential disruptiveness of AWS in the conduct of war that objections from *ad bellum* have been raised against their use (Asaro 2008; Roff 2015). The objections that we consider here are of two sorts. Both relate to the possibility that AWS may increase the incidence of war: the first is that, by reducing the costs—broadly conceived, e.g. as both human and economic—for going to war, AWS will reduce the threshold for going to war. Second, and relatedly, providing a propagandistic value, AWS will embolden political leaders to lead an otherwise unwilling population into support for war.

Each of these objections will be considered in turn under the headings of proportionality and last resort. The argument made in this article is twofold. The first is that these objections—particularly, the objection from proportionality—are not developed with adequate precision to stand, as they are, as objections to the use of AWS. The second is that, whilst these objections present pressing concerns in their own right, demanding thoughtful consideration, they are not problems for Just War theory *ad bellum* principles per se. Whilst *jus ad bellum* provides a powerful set of principles for assessing the justifications for war; objections based on these principles against AWS are not conceptually compelling, for they fail to account for the changing nature of war, and for the fact that ethical considerations on war must ‘track’ this transformation. Encapsulating this sentiment, Abney has noted that,

“...new capabilities transform not only the conduct of war, but also the very understanding of what war is, and when and how it ought (not) to be waged. Accordingly, such innovations require clarifications, if not wholesale revisions, to ethical concepts and theories.” (Keith Abney 2013, 338)

To be sure, in seeking to refute arguments made thus far from proportionality and last resort, we are not arguing that because *ad bellum* objections do not hold conceptually then AWS are acceptable or justifiable, nor that in practice, the scenarios envisaged by the two sets of objections will never materialize. We mean here only that *jus ad bellum* does not provide the right conceptual tool to assess the justifiability or ethical legitimacy of AWS. Our hope is that in providing these clarificatory remarks, the debate over the ethical uses of AWS can be made more rigorous.

2 Ad bellum proportionality

For a war to be considered ‘just’, there must be an overall proportionate relationship between the destruction caused by the war and the good the war will do (Hurka 2005, 35).

² For sake of simplicity we refer to Coverdale (2004) in listing the Just War principles. There is some divergence over the exact number of principles pertaining to Just War theory amongst Just War theorists. For instance, some theorists include the eschewal of ‘means mala in se’ in the list of *in bello* principles. However, this divergence is more to with the porous character of Just War theory rather than fundamental disagreement over the nature of those principles.

³ The strict separation of *jus ad bellum* and *jus in bello* has been contested by revisionist Just War theorists, most influentially by Jeff McMahan (see: Jeff McMahan 2009). Here there is not the scope to rehearse this debate; in this article we consider Just War theory in so far as it relates to IHL.

Weighing up the overall costs and benefits of war requires a global outlook (what are the costs to non-combatants on all sides) and it entails conceiving of ‘costs’ in terms of both material and moral costs. The *ad bellum* proportionality calculus is difficult, because it is entirely future-orientated, including a number of indeterminacies and unknown quantities. In weighing up overall whether a war is proportionate, estimates must include the number of lives that would have been lost, as well as the structural and economic damage caused by waging the war. This is no small feat given the unexpected and unintended consequences often associated with armed conflict. *Ad bellum* proportionality, as Meisels argues “always requires an essentially inaccurate prediction of prospective scenarios and foreseeable (or unforeseeable) danger [...]” (Tamar Meisels 2018, 76). This is why, while recognising its theoretical relevance to the Just War tradition, the difficulties in implementing the *ad bellum* proportionality principle to concrete cases lead Walzer to claim it to be a relatively minor restriction on war (Michael Walzer 2009; see also: Forge 2009). In this article, we are not concerned with assessing the relevancy of *ad bellum* proportionality, and Walzer’s conclusion. We share Walzer’s concerns and focus on the challenges of applying this principle to the case of AWS.

Despite Walzer’s claim about its ‘minor’ status as an *ad bellum* restriction on war, proportionality has loomed large in objections to AWS. The objection here is that AWS will reduce the costs of war and thereby weight the overall proportionality calculus towards its benefits for the side which deploys them. Thus, promising more proportionate wars, AWS are likely to make war more likely and increase the overall global incidence of war. As Abney writes,

“autonomous robots, with their promise of fewer casualties, will make war less terrible and therefore more tempting, plausibly enticing political leaders to wage war more readily” (Keith Abney 2013, 340).

These objections chime with perennial concerns about the development and deployment of technology to ‘humanize’ war as fundamentally misconceived, because by rendering war more tolerable, this technology makes war more likely. As Clark writes,

“It has long been held that humane warfare, on this reasoning—whatever the good intent underlying it—is subversive of the goal of reducing the incidence of war, and so finally eliminating it” (Ian Clark 2015, 118).

There are three main limitations to this objection: overlooking the difficulties of calculating *ad bellum* proportionality; confusing the concept of proportionality of effects with the precision of weapon systems; disregarding the ever-changing nature of war and of its ethical

implications. In the rest of this section, we will analyse each limitation in turn.

Let us consider first the *ad bellum* assessment of proportionality. On one hand, there are substantial methodological difficulties involved in quantifying the costs of war, particularly in making discrete sets of values commensurate. The precision offered by AWS may reduce costs in terms of lives lost, but there are concerns that the algorithmic decision to kill without human control infringes the principle of human dignity (Heyns 2017). Thus, the costs diminished by the precision of the weapon could be offset by moral costs, if the use of AWS is found to infringe the principle of dignity in death (Horowitz 2016). On the other hand, is the difficulty in estimating the net effects of a given technology. As Sechser et al. write:

“Extrapolating from current technological trends is problematic, both because technologies often do not live up to their promise, and because technologies often have countervailing or condition effect that can temper their negative consequences.” (Sechser et al. 2019, 728)

This leads to the second limitation, that for each war fought in a future series of wars, we would have to calculate how the severity and duration of each of those wars fought affects the severity and duration of the next. For this, there is no reliable calculation because there is no rule for doing so. Walzer has made this point, in reply to the utilitarian argument of General von Moltke that restrained warfare prolongs fighting, whilst “the greatest kindness in war is to bring it to a speedy conclusion.” Walzer replies.

“But if we imagine a series of wars, this argument probably won’t work. At any given level of restraint, let’s say, a war will take so many months. If one of the belligerents breaks the rules, it might end more quickly, but only if the other side fails or is unable to reciprocate. If both sides fight at a lower level of restraint, the war may be shorter or longer; *there isn’t going to be any general rule.* (Michael Walzer 1977, 131, emphasis added).”

Ad bellum proportionality calculations are difficult and indeterminate enough with wars forecast for the near-future. Including in that calculus wars that not only are yet to happen but also are unforeseeable generates a high degree of indeterminacy which undermines arguments about the likelihood of incidence of war. These difficulties are compounded when trying to relate *in bello* proportionality calculations to *ad bellum* proportionality calculations, because they relate to two areas which are logically distinct in Just War theory.

The argument against AWS based on the principle of proportionality is premised on a common-sense understanding of proportionality rather than a rigorous application of

Just War principles, and because of this, it trivialises the complexity of proportionality calculations and overlooks the dynamic nature of war. It is important to remark here that even if weapons lend themselves to uses which allow greater (or lesser) respect for the principle of proportionality, proportionality is not something which inheres to weapons themselves.

This misconception about the nature of proportionality is common enough that it forms the basis of arguments made by both those arguing for and against AWS. Both views—those advocating for AWS under the assumption that they lead to ‘more proportionate’ war waging and those who argue for a ban of AWS—confuse ‘proportionality’, a contextual assessment, with ‘precision’, an objective property of weapons (Braun and Brunstetter 2013).⁴ AWS may well be more precise in terms of fixed and objective measures such as blast radius and other technical attributes; but this is irrelevant for the *ad bellum* proportionality calculus. This has to do with the assessment of overall purposes sought in the pursuit of a war, the tactics, the types of costs envisioned, etc.

The misconception about the nature of proportionality is rooted in a naïve understanding of the idea of war, the third limitation of the objection. Both those who argue for and against AWS on the grounds of their proportionality (or lack of) in this way assume an unchanging idea of ‘war’ along which fluctuating incidences of war can be modelled and upon which a static ethics of war—including proportionality considerations—can be based. Instead, it is important to note that technologies of war themselves generate transformations in the nature of war and, thus, the ethics of war. This point is key to the argument we are making here. For centuries, debates about the ethics of war have been driven by technological innovations: crossbows, gunpowder, submarines, aerial bombardment, nuclear weapons, lasers, drones, and now, the use of artificial intelligence in war. Technological innovation has always led to a reevaluation not just of how the principles of Just War are to be applied, but also of how they are interpreted. As Clark has argued, technological innovation does not generate straightforwardly new ethical problems within the boundaries of an unchanging idea of war; rather, new ethical problems arise because of the way that

technological innovation first impacts our understanding of war, and consequently, normative analyses and approaches to its regulation. (Ian Clark 2015, chap. 7 War, Technology, and Conceptual Change).

Nuclear weapons and drones are two technologies that provide a stark illustration of this point. The advent of nuclear weapons, which obliterated the categories of discrimination, was said to take war “past a boundary at which many previous concepts and categories of appraisal—both military and political—ceased to apply, or even to have meaning” (Quinlan cited in: Lawrence Freedman 2018, 92). When Walzer declared that nuclear weapons “explode the theory of just war” (Walzer 1977, 282), it was because they upended the idea of war upon which Just War theory had until then been premised, one wherein the possibility of discrimination at least made conceptual sense.

Similarly, the debate over the introduction of drones was spurred for the fact that they challenged the age-old conception of war as, in essence, entailing a ‘contest’ comprised of (at least) two sets of opposing combatants experiencing mutual risk (Steinhoff 2013; Braun and Brunstetter 2013; Bradley Jay Strawser 2013; Schulzke 2016). Of course, warfare has never entailed complete symmetry of risk, but commentators have argued that, since drones entail in essence the complete avoidance of risk for the combatant, they cannot be plausibly reconciled with the traditional conception of war as a ‘contest’. This is ethically problematic, because if drones represent an instrument of violence irreconcilable with the nature of war, then the ethical principles designed for war—*jus ad bellum*, *jus in bello*—are inapplicable to the use of drones. If, as Enemark writes, “drones resemble ‘a godlike power to call down destruction from the skies,’” the rules for restraining such strikes would need to be derived from a different concept of violence” (Enemark 2014, 368). Work on drones has henceforth sought to develop an adequate conceptual basis outside of ‘war’ for “restraining the resort to this unfamiliar form of violence” (Enemark 2014, 370).

The point here is that specific ethical debates about war make sense only once prior conceptual framings of the war being waged are understood properly. “As a result, when we apply ethics to war, we are left to shoot at a constantly moving target, and the ethics have to track that evolution” (Ian Clark 2015, 19). The objection that AWS will increase the incidence of war forgets this moving target and treats war as an unchanging phenomenon. Pre- and post-AWS war cannot be compared as like for like, as AWS are a transformative innovation that alters the way in which war may be fought and the measures by which the overall incidence can be assessed. That is not to say that fears about the rush to war are misplaced, it is only to intercede with a prior question: “rush to what?”.

⁴ The confusion generated by overlooking such factors has a long pedigree and was evidenced quite starkly in the initial deployment of drones in combat settings and the argument, made by their proponents, that, once viewed on an “historical trend line” drones can be understood as generating more proportionate wars (Anderson n.d., 383–84). However, often this ‘historical trend line’ took the second world war as its base line, and given that the allied forces intentionally targeted civilian population, it seemed not much praise to model drones against a “historic nadir for warfare” (Braun and Brunstetter 2013, 309; See also: Shue 2008).

3 Last resort

The second objection to the use of AWS on which we focus in this article is centred on the principle of last resort.⁵ According to this objection, reduced costs would make the choice to go to war politically and economically more convenient and thus state actors may decide to declare war instead of pursuing alternative means to resolve a conflict, thus violating the principle of last resort. Sharkey raises the image of the ‘body-bag count’ as one of the greatest ‘inhibitors’ of military action: fewer body-bags “means fewer disincentives to start wars” (Sharkey 2008, 16). The concern that political leaders have for “body-bags,” as Sharkey writes, is the concern for the “body-bags” returning home, not the number of fatalities in a given war overall—the fatalities of enemy combatants, of non-combatants, and of overall material damage, follow as only a secondary consideration (Sharkey 2008, 16; Sharkey 2012, 2016). In this way, AWS will embolden political leaders to undertake wars without exhausting less destructive courses of action.

Asaro takes this concern to have particular relevance to democratic states where there is a premium on the propaganda that political leaders are able to muster to lead an otherwise unwilling public to support war. As he writes,

“Irrespective of the underlying justness of the motives, when the leadership of a state decides to go to war, there is a significant propaganda effort. This effort is of particular importance when it is a democratic nation, and its citizens disagree about whether a war is worth fighting and there are significant political costs to a leader for going against popular sentiments. A central element of war propaganda is the estimation of the cost of war in terms of the lives of its citizens, even if that is limited to soldiers, and even if those soldiers are volunteers. A political strategy has evolved in response to this, which is to limit military involvement to relatively ‘safe’ forms of fighting in order to limit casualties, and to invest in technologies that promise to lower the risks and increase the lethal effectiveness of their military” (Asaro 2008, 8).

Asaro thereby concludes that.

“one of the strongest moral aversions to the development of robotic soldiers stems from the fear that they will make it easier for leaders to take an unwilling nation into war” (Asaro 2008, 7).

The principle of last resort represents the preference in the Just War tradition for peace over war. It states that political leaders are obligated to assess all available means for meeting a given threat and to opt for the means sufficient for doing so given this preference for means other than war (Coverdale 2004, 258–59). Put in the negative, war is to be the “option least to be preferred”. Or, put differently still, war is to be “as late as possible, as early as necessary” (Quinlan 1997, 16).

Thus, no matter if AWS provide an incentive—or propaganda value—in going to war, if other means have not been exhausted properly then the decision to resort to war remains unjust. As such, whilst Asaro discusses the above objection under the rubric of Just War theory, the breaching of the principle of last resort is less inherent to AWS and more linked to the willingness of political leaders to abide by this principle. That is an empirical question about whether political leaders are prepared to abide by the principles of Just War theory, not about the practicability of Just War principles themselves. There have been attempts to marry psychology and cognitive science with Just War theory to explain collective self-deception in going to war (Richard Werner 2013), but explaining why states fail to live up to the demands of Just War theory does not tell us what, in this context, those demands are.

However, there is a point to be raised about the influence of technology when it comes to making judgements about last resort. It is worth remembering that the principle of ‘last resort’ does not mean that nations should turn to war when no other possible course of action is available. Interpreting last resort in this way would mean that war can never be justified since it is always possible to say that not every alternative has been tried. Rather, the principle of last resort.

“requires a considered judgement about whether some imagined alternative has a good chance of avoiding war. It does not require that every idea actually be pursued to the end of the line” (Allen quote in: Coverdale 2004, 259)

More sophisticated work on technology in Just War theory considers not only the ethical problem that instruments of warfare pose, but how such instruments play a part in forming or influencing judgements to do with proportionality or last resort. In which case, the objection discussed in this section has some value, insofar as new technological developments particularly around human–machine teaming (Ministry of Defence 2018) require vigilance to identify unpredicted and unwanted effects, like for example escalation. As Allenby writes,

“...a focus only on the physical technologies themselves will be entirely inadequate to consideration of deeper questions of technological impact... it is

⁵ Much of the opposition to AWS on the basis of last resort parallels concerns that have been raised over the use of drones in warfare (see: Bradley J. Strawser and McMahan 2013). For sake of clarity we maintain a focus here on AWS.

impossible to understand the implications of planned or potential military and security technologies without a deeper initial understanding of the cultural, operational, and institutional frameworks within which the technologies are being conceived, and will be used” (Allenby 2013, 293).

An example of this is the reliance of western military on the so-called ‘technological third-wave’, increasingly placing digital information processes at the centre of organisations (Lawrence Freedman 2018, 185–86). The belief is that increased information processing and communication would generate greater precision in weapons use and offer swift victory through the identification of enemy vulnerabilities in complex settings while limiting risk to one’s own troops. As Freedman writes, by the mid-1990s,

“a vision of war was developing which would get the whole affair over quickly with few casualties. Extracting the pain from war was essential to the project. If war could become both high-impact and low-casualty, then it could be socially contained and retained as a political instrument” (Lawrence Freedman 2018, 189).

Given the complexity and destructiveness of war there is little surprise that technological fixes, promising to turn violence into a manageable and containable instrument of politics, generate much interest, not least if those fixes can attenuate the political costs of war by reducing the casualties on one’s own ‘side’. However, whilst such hopes for “Dominant Battlespace Knowledge” and “Near-Perfect Mission Assignment” might be realizable in battlespaces largely empty of combatants, such as the ocean, they have proven to be quite problematic when such technologies were confronted with complex scenarios, like the type of urban insurgencies western forces were to meet in Iraq (Freedman 2018, 157–97). The belief that casualties, through technological fixes, *could* be negligible, if not zero, morphed into an *ought* to be negligible. This, as Walzer (2004, 99–103) argues, engendered the type of warfare seen in Kosovo where high-altitude bombing, meant to reduce the risks to NATO forces, weighted the risks of war unacceptably towards civilians. For our current discussion, it is noteworthy that such a conflict—Kosovo—fought in such a way as to minimize NATO casualties, is also a conflict which has been criticised using Just War theory for failing to live up to the *ad bellum* criteria of last resort (Robinson 1999).

4 Conclusion

Whilst the type of *ad bellum* objections raised against AWS explored here are conceptually unsound, they do have the merit of pointing to the way that instruments of war can

contribute to the proliferation and escalation of war under the influence of myths about technological innovation and precision violence. However, such objections, which have an intuitive plausibility, risk being misleading in assessing the ethical impact of AWS and deciding upon, and regulating, their use. The limits of the *ad bellum* objections analysed in this article do not entail that AWS are acceptable or justifiable, nor that in practice the scenarios envisaged by the two objections will never materialize. The analysis of these limits offers two contributions to the debate on AWS. First, our analysis has made clear that *ad bellum* principles are not the best set of ethical principles for tackling the ethical problems raised by AWS. Second, to ask the right questions about the ethical problems raised by AWS, we must first track the way that AWS transforms the nature of war itself.

Funding Alexander Blanchard and Mariarosaria Taddeo have been funded by the Dstl Ethics Fellowship held at the Alan Turing Institute. The research underpinning this work was funded by the UK Defence Chief Scientific Advisor’s Science and Technology Portfolio, through the Dstl Autonomy Programme, grant number R-DST-TFS/D026. This paper is an overview of UK Ministry of Defence (MOD) sponsored research and is released for informational purposes only. The contents of this paper should not be interpreted as representing the views of the UK MOD, nor should it be assumed that they reflect any current or future UK MOD policy. The information contained in this paper cannot supersede any statutory or contractual requirements or liabilities and is offered without prejudice or commitment.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abney K (2013) Autonomous robots and the future of just war theory. In Routledge handbook of ethics and war: just war theory in the twentieth-first century. London
- Allenby B (2013) Emerging technologies and just war theory. In: Routledge handbook of ethics and war: just war theory in the twenty-first century. London: Routledge.
- Amoroso D, Guglielmo T (2020) Autonomous weapons systems and meaningful human control: ethical and legal issues. *Curr Robot Rep* 1(4):187–194. <https://doi.org/10.1007/s43154-020-00024-3>
- Anderson K (2021) Efficiency in bello and ad bellum: making the use of force too easy? Targeted killings. Oxford University Press, Oxford, Oxford
- Asaro P (2008) ‘How just could a robot war be’. In: Current issues in computing and philosophy, pp 50–64.

- Birnbacher D (2016) Are autonomous weapons systems a threat to human dignity? In: Claus K, Hin-Yan L, Nehal B, Robin G, Susanne B (eds) *Autonomous weapons systems: law, ethics, policy*, 105–21. Cambridge University Press, Cambridge. <https://doi.org/10.1017/CBO9781316597873.005>
- Braun M, Brunstetter DR (2013) Rethinking the criterion for assessing CIA-targeted killings: drones, proportionality and jus ad vim. *J Milit Ethics* 12(4):304–324
- Chengeta T (2016) Measuring autonomous weapon systems against international humanitarian law rules. *J Law Cyber Warfare* 5(1):66–146
- Clark I (2015) *Waging war: a new philosophical introduction*. Oxford University Press, Oxford
- Convention on Certain Conventional Weapons. 2019. ‘Report of the 2019 Session of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems’. CCW/GGE.1/2019/3. Geneva: United Nations Office for Disarmament Affairs. <https://undocs.org/pdf?symbol=en/CCW/GGE.1/2019/3>.
- Coverdale JF (2004) An Introduction to the Just War Tradition. *Pace International Law Review* 16(2):221–278
- Crootof R (2016) War torts: accountability for autonomous weapons. *Univ Pa Law Rev* 164(6):1347–1402
- Demy TJ (2020) “‘Something old, something new’—reflections on AI and the just war tradition”. In: Masakowski YR (ed) *Artificial intelligence and global security*. Emerald Publishing Limited, pp 53–62. <https://doi.org/10.1108/978-1-78973-811-720201003>.
- Enemark C (2014) Drones, risk, and perpetual force. *Ethics Int Aff* 28(3):365–381. <https://doi.org/10.1017/S0892679414000446>
- Finlay C (2019) *Is just war possible?* Polity, Cambridge
- Forge J (2009) Proportionality, just war theory and weapons innovation. *Sci Eng Ethics* 15(1):25–38. <https://doi.org/10.1007/s11948-008-9088-z>
- Foy J (2014) *Autonomous Weapons Systems: Taking the Human out of International Humanitarian Law*. *Dalhous J Leg Stud* 23:47–70
- Freedman L (2018) *The future of war: a history*. Penguin
- Grut C (2013) The challenge of autonomous lethal robotics to international humanitarian law. *J Conflict Secur Law* 18(1):5–23
- Heyns C (2017) Autonomous weapons in armed conflict and the right to a dignified life: an African perspective. *S Afr J Hum Rights* 33(1):46–71
- Himmelreich J (2019) Responsibility for killer robots. *Ethical Theory Moral Pract* 22(3):731–747. <https://doi.org/10.1007/s10677-019-10007-9>
- Horowitz MC (2016) The ethics and morality of robotic warfare: assessing the debate over autonomous weapons. *Daedalus* 145(4):25–36
- Hurka T (2005) Proportionality in the morality of war. *Philos Public Aff* 33(1):34–66
- Jain N (2016) Autonomous weapons systems: new frameworks for individual responsibility. In: Bhuta N, Beck S, Geiß R, Liu H-Y, Kreß C (eds) *Autonomous weapons systems: law, ethics, policy*. Cambridge University Press, Cambridge
- McMahan J (2009) *Killing in war*. Oxford University Press, Oxford
- Meisels T (2018) *Contemporary just war: theory and practice*. Routledge, London
- Ministry of Defence (2018) *Human–machine teaming (JCN 1/18)*. <https://www.gov.uk/government/publications/human-machine-teaming-jcn-118>.
- Quinlan M (1997) ‘Thinking about nuclear weapons’. Whitehall Paper. London: Royal United Services Institute for Defence Studies.
- Roach SC, Amy EE (2020) *Moral responsibility in twenty-first-century warfare: just war theory and the ethical challenges of autonomous weapons systems*. SUNY Press.
- Robinson P (1999) “Ready to kill but not to die”: nato strategy in Kosovo. *Int J. Autumn* 12
- Roff HM (2015) Lethal autonomous weapons and jus ad bellum proportionality. *Case Western Res J Int Law* 17
- Sassoli M (2014) Autonomous weapons and international humanitarian law: advantages, open technical questions and legal issues to be clarified. *Int Law Stud Naval War Coll* 90:308–340
- Schulzke M (2013) Autonomous weapons and distributed responsibility. *Philos Technol* 26(2):203–219. <https://doi.org/10.1007/s13347-012-0089-0>
- Schulzke M (2016) The morality of remote warfare: against the asymmetry objection to remote weaponry. *Polit Stud* 64(1):90–105. <https://doi.org/10.1111/1467-9248.12155>
- Sechser TS, Narang N, Talmadge C (2019) Emerging technologies and strategic stability in peacetime, crisis, and war. *J Strateg Stud* 42(6):727–735. <https://doi.org/10.1080/01402390.2019.1626725>
- Sharkey N (2008) Cassandra or false prophet of doom: AI robots and war. *IEEE Intell Syst* 23(4):14–17
- Sharkey N (2010) Saying “No!” to lethal autonomous targeting. *J Milit Ethics* 9(4):369–383. <https://doi.org/10.1080/15027570.2010.537903>
- Sharkey NE (2012) The evitability of autonomous robot warfare. *Int Rev Red Cross* 94(886):787–799
- Sharkey A (2019) Autonomous weapons systems, killer robots and human dignity. *Ethics Inf Technol* 21(2):75–87. <https://doi.org/10.1007/s10676-018-9494-0>
- Sharkey N (2016) Staying in the loop: human supervisory control of weapons. In: Claus K, Hin-Yan L, Nehal B, Robin G, Susanne B (eds) *Autonomous weapons systems: law, ethics, policy*. Cambridge University Press, Cambridge, pp 23–38
- Shue H (2008) Concept wars. *Survival* 50(2):185–192
- Sparrow R (2007) Killer robots. *J Appl Philos* 24(1):62–77
- Steinhoff U (2013) Killing them safely: extreme asymmetry and its discontents. In: Bradley JS (ed) *Killing by remote control: the ethics of an unmanned military*. Oxford University Press. <https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780199926121.001.0001/acprof-9780199926121-chapter-9>.
- Strawser BJ, McMahan J (eds) (2013) *Killing by remote control: the ethics of an unmanned military*. Oxford University Press, Oxford
- Strawser BJ (2013) Introduction: the moral landscape of unmanned weapons. In: Bradley JS (ed) *Killing by remote control: the ethics of an unmanned military*. Oxford University Press, Oxford, pp 3–24
- Taddeo M, Blanchard A (2021) A comparative analysis of the definitions of autonomous weapons. SSRN 3941214.
- Taddeo M, Blanchard A (2022) Forthcoming. ‘Ascribing moral responsibility for the actions of autonomous weapons systems (Forthcoming)’.
- Taddeo M, McNeish D, Blanchard A, Edgar E (2021) Ethical principles for artificial intelligence in national defence. *Philos Technol* 34(4):1707–1729. <https://doi.org/10.1007/s13347-021-00482-3>
- US Department of Defense (2012) ‘DoD Directive 3000.09 on Autonomy in Weapon Systems’. <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>.
- Walzer M (1977) *Just and unjust wars: a moral argument with historical illustrations*. Basic Books, New York
- Walzer M (2004) *Arguing about war*. Yale University Press, London
- Walzer M (2009) ‘On Proportionality’. *The New Republic*, 8 Jan 2009. <https://newrepublic.com/article/64580/proportionality>.
- Werner R (2013) Just war theory: going to war and collective self-deception. In: Allhoff F, Evans NG, Henschke A (eds) *Routledge handbook of ethics and war: just war theory in the twenty-first century*. Routledge, London