

Drones, counterfactuals, and equilibria: Challenges in evaluating new military technologies

Toby Ord*

Using drones as my primary example, I explore some of the challenges in evaluating new military technologies. I show how we need a deep understanding of the relevant counterfactuals in order to make accurate and nuanced evaluations of military technologies. I then explore the force of the argument that introducing a new military technology can lower the cost of going to war, finding that it has very wide reaching effects that go far beyond the case of drones. I conclude that it is exceptionally difficult to accurately evaluate new military technologies and that we need to learn new skills and techniques to deal with the empirical challenges involved.

The use of armed drones in war has raised significant moral concern. At the time of writing, at least two dozen countries have operational armed drones in their militaries. In Pakistan, US drones have already killed more than two thousand people, more than ten percent of whom were civilians.¹ Both in academia and in the wider public discourse, many people have suggested that the use of these drones is wrong, that countries should forgo using them, and that treaties should be enacted to ban them.² Less commonly, there have been attempts to defend their use and even argue that we are morally required to use them where possible.³

Drones are technically known as unmanned aerial vehicles (UAVs). While they have many other uses in both civilian and military life (such as surveillance and delivery), the focus of this paper is on weapon-carrying drones deployed by the military, such as the MQ-1 Predator and MQ-9 Reaper deployed by the US in Afghanistan. When I use the term ‘drone’ in this paper I only mean to refer to military UAVs capable of delivering deadly force, and I only mean to discuss their use in war. All such drones currently require a human on the ground to authorise the use of their weapon systems (a so-called *human in the loop*), though in the future we may see drones that are able to select their own targets. Much of what I say would also apply to such

* Oxford Martin Programme on the Impacts of Future Technology, Faculty of Philosophy, University of Oxford. [[Acknowledgements to go here](#)].

¹ (Serle & Ross, 2014) estimates 2,412–3,701 total deaths, with 416–951 civilian deaths; (New America Foundation, 2014) estimates 2,080–3,428 in total, of which 258–307 are civilians. These numbers may significantly underestimate the number of civilians killed due to methodological problems, see (Grut *et al.*, 2012).

² See, for instance, (O’Connell, 2012) and (Human Rights Watch & International Human Rights Clinic, 2012). For a discussion of the arguments, see (Whetham, 2013).

³ See, for instance, (Strawser 2010).

fully autonomous systems, but my focus here will be on those with a human in the loop.

In this paper, I discuss two important tools for thinking about the ethics of introducing new military technologies: counterfactuals and equilibria. I use drones as the central example of a new military technology, but the points I make apply much more broadly. I argue that without a careful understanding of these tools of thought, the ethical analysis of such new technologies is greatly impoverished and cannot be relied upon. I don't claim to have invented these tools: one is a mainstay of academic ethics and the other is a central concept in economics. However, I shall show that we need a much better understanding of them and how they relate to the ethics of war in order to make any real progress in the moral evaluation of new military technologies.

I will focus on evaluating the *consequences* of introducing a new military technology. This is not the only type of ethical question one could ask. For instance, I won't address whether it is *honourable* to kill in war without exposing yourself to risk of physical harm, or whether there might be ethical *side-constraints* which prohibit lethal use of drones regardless of the benefits. However, questions about the moral consequences of deploying a new technology are central to much of the discussion and include such key questions as 'Will the deployment of drones lead to more civilian deaths in war?' and 'Would the effects of unilateral relinquishment of drones by the US be good for the world?'. I presume all parties — be they consequentialists, deontologists, or virtue theorists — are interested in such questions.⁴

Counterfactuals

Drones are high-technology devices created for the purpose of killing people and destroying property. As well as killing enemy combatants, drones have killed hundreds of civilians. In the near future, many more drones will be built and deployed by dozens of countries and will kill innocent people across the globe. Facts like these might lead people to conclude that drones are bad, that their deployment is wrong, and that we should take stands to oppose their use.⁵

This is too quick. We cannot legitimately draw such a conclusion from these facts alone. Of course drones are destructive and deadly — they are a military technology. The same is true of swords, bows, catapults, guns, artillery, and bombers. Like all other weapons of war, drones will kill civilians in some situations and occasionally this will even be deliberate. The real questions are 'How many?' and 'How often?'. Do they lead to more civilian killings than aerial bombing campaigns? Will their widespread introduction raise the number of civilian killings or lower it?

⁴ As John Rawls put it: 'All ethical doctrines worth our attention take consequences into account in judging rightness. One which did not would be irrational, crazy.' (Rawls 1971, 30).

⁵ For instance John Forge argues that researching new weapon technologies is wrong because it provides others with the means to harm people without appropriate justification (Forge 2013, 24–36, 135–53).

We cannot assess drones in isolation. Instead, we need to pay careful attention to what they will replace and compare the two alternatives. We need to consider the *counterfactual*: what would happen if drones were not deployed.

When a new technology is superior to an old one that occupied largely the same niche, the old technology gets used less. If the new technology is sufficiently superior, the old technology becomes obsolete and ceases to be used at all. This is as true in the military as in everyday life. Cannons crowded out catapults, and were themselves crowded out by modern artillery. Muskets crowded out bows, rifles crowded out muskets, and automatic weapons have largely crowded out rifles.

What will drones replace? There is no single answer. Even with our focus on armed drones, there are many objectives that drones can be put towards and many existing technologies that were previously used to achieve these objectives.

Technologically, most drones resemble aeroplanes. They are aerial platforms from which the user can deploy weapons to kill combatants and destroy property on the ground with limited possibility for retaliatory strikes. This has significant overlap with the role of bombers: aircraft that primarily drop bombs and missiles at targets on the ground. In order to avoid detection and retaliation, bombers are often flown at high altitudes and sometimes used at night. This makes it harder for them to accurately hit their targets, increasing the damage to civilian property and life. Bombers were famously used in the attacks on London and Dresden in the Second World War, the carpet bombing of the Vietnam War, and the aerial bombardment of Kosovo.

While drones are not completely reliable in their targeting, they may well be significantly better than bombers. Their lower financial cost and lack of an onboard pilot mean that more risks can be taken with them, including flying closer to their targets. This provides a better view of whether it is indeed a valid military target and can lower the chance of missing the target — hitting something else instead.

So despite being destructive weapons, the introduction of drones into warfare has the potential to lead to *less* damage to civilian lives and property. They are quite plausibly better than bombers, and to the extent to which they replace bombers, they may make modern warfare less damaging to civilians.

People who reject the use of drones should have no love of bombing. The fire bombing of Dresden and Tokyo, and the carpet bombing of Vietnam, were among the most horrible episodes of modern warfare. So if more drones means less bombing, that may be a very good thing. Among other things, it would mean that attempts to develop a treaty to ban the use of drones or to encourage a country to unilaterally relinquish the use of drones may predictably cause more civilians to be killed in war.

On the other hand, much of the current use of drones occurs in Pakistan where they are primarily used to target particular, high-value combatants. It is not clear whether this is indicative of how they will be used in the future. This targeted killing is traditionally the role of troops, police, or covert operations on the ground rather than

bombers. Here the ability to accurately determine whether or not the person is a valid military target, and to accurately deploy force if they are, is plausibly higher for the ground troops than for the drones.⁶ If so, then to the extent to which drones replace ground forces, they may make modern warfare *more* damaging to civilians.

Whether the deployment of drones in war has good or bad effects overall thus depends on the counterfactual situation. To what degree is it crowding out bombing versus on the ground forces, or something else? This is a key empirical question, though not one that is often asked by opponents (or proponents) of the technology. Moreover, its answer will depend on exactly what form of question is being asked.

For instance, consider the questions ‘Should the US cease deploying drones in war?’ and ‘Should we have an international treaty to prohibit the use of drones in war?’. It may well be that the ratios of replaced activities will be different between the US and the world as a whole, which (among other things) may make these questions have different answers.

As well as these overall questions of whether the use of drones would be good or bad given the empirical facts about what will be crowded out, consideration of counterfactuals can also lead us to make more nuanced points. For example, we may decide that it is good to use drones for those missions that were previously done by bombers, and bad to use them for those missions that were previously done by forces on the ground. We may then try to encourage the first of these while discouraging the second. Perhaps the first could be encouraged by trying to establish a treaty to ban aerial bombing rather than banning drones, or to get the US to unilaterally cease aerial bombing. The second could be discouraged by trying to establish a treaty or an internal military doctrine to prevent the use of drones in targeted killing, rather than preventing all drone use.

The importance of carefully considering the relevant counterfactual is also crucial when considering non-technological changes in how we conduct war. As just one example, take the ‘decapitation’ strategy of attempting to kill the enemy’s military and political leaders, as used in Kosovo and the Iraq War. This has been advocated by historical figures such as Sun Tzu,⁷ but also receives significant public condemnation.⁸ The best argument for such a strategy probably lies in considering the counterfactual. Killing political leaders in their beds is horrible, but so is the day to day activity of war in which thousands or millions are killed and many more atrocities committed. I don’t know if decapitation strikes are justified, but if they are it is presumably because while terrible, the alternatives are even worse.

⁶ Note that preliminary evidence actually suggests the opposite. Avery Plaw’s compilation of figures from Pakistan from 2004–10 states that drones have had a ratio of intended military targets to civilian deaths of 17 to 1, while Pakistan’s special forces and army have only had ratios of 4 to 1 and 3 to 1 respectively (Plaw 2010). See also (Strawser 2010, 351–2).

⁷ (Tzu 1963).

⁸ For instance, (Wright 2010); (Russell 2013). See also (Brough & Williams, 2011) for a good introduction; and for more analysis (Jordan 2009) and (Johnston 2012).

Equilibria

A major argument against the widespread deployment of drones in war is that by avoiding risking the lives of their pilots, they would lower the cost of going to war. This would then increase the number of wars that drone-deploying countries enter, making the use of drones very bad overall.⁹

The basic point here is a very important one and fundamental to economics. We often think about the direct effect of actions, new technologies, and policies, but they can also have systematic indirect effects. In economics, an *equilibrium* is a state where economic forces such as supply and demand are balanced. In our case, before the introduction of drones the national willingness to go to war and the number of wars entered were in equilibrium. By lowering the cost of going to war, the deployment of drones upsets one of these forces which causes the other to adjust until they reach a new balance. In this case the new equilibrium will be one where more war will be 'consumed' by the deploying country. Given that more war is typically bad, this effect of changing the equilibrium could potentially be large enough to outweigh any direct ethical benefits that deployment of drones might create.

The deployment of drones *will* lower the cost of going to war. For a given set of military objectives, using drones will mean fewer casualties on the deploying side and quite likely less monetary cost too, especially when reduced combat training and reduced military pensions are factored in. Reduced costs of going to war will indeed increase one's country's appetite for it. It would be difficult to know how many more wars this will cause that country to start, but if they are known to have fewer costs, then the public and the government will be less reluctant to start or join a conflict and, statistically speaking, we should expect more wars.

However, in this simple form, the argument has a number of lacunae. Firstly, while we should expect the number of wars to increase, this is because the costs of war to the deploying side in terms of human life (and possibly money) go down, which is itself a moral benefit. It is not clear what happens to the badness of war overall. Perhaps the increase in the number of wars is more than offset by the reduced loss of life in them, making the situation less bad all things considered. It is very difficult to know. This may give reason to change the conclusion of the argument to say only that the increased number of expected wars in the future makes drones worse than they would otherwise appear — possibly bad enough to outweigh their other benefits, though this isn't clear.

A second problem concerns wars that would (all things considered) be good to start or to join. Examples might include joining the Second World War on the side of the Allies, or starting an intervention in Rwanda to stop the genocide. When it comes to such wars, reducing the cost of going to war (and thus increasing the appetite for war) might actually be a good thing, making it easier for countries to intervene to help others. I do not know if there are times it is right to start or join a war, but those

⁹ For a critical discussion of this argument, see (Beauchamp & Savulescu 2013). See also (Strawser 2010, 358–61).

who think they do exist and are not extremely rare might find that this makes reducing the cost of going to war less bad.

A third and more subtle problem arises when you consider that all new military technologies lower the cost of going to war for the country that first deploys them. Otherwise they would choose not to deploy them. Thus the introduction of catapults, cannons, artillery, and so forth all decreased the cost of going to war for the deploying sides. This would have initially made the deploying sides slightly more willing to go to war, but after other sides deployed these technologies too, the overall costs of going to war may have gone back to normal or even increased. In the case of new methods of bombarding cities such as those above, I presume that the costs did increase overall. In the case of drones, it is not clear what happens to the cost of going to war once all parties have access to them. Thus the period of reduced costs of going to war may just be temporary. This wouldn't make the argument fail, but would remove some of the sting.

Overall, I think the argument that the deployment of drones would lead to more war retains significant force. It shows that drones have a hidden moral cost and suggests (though does not prove) that this may even be large enough to offset their advantages in protecting their operators from harm. However, if this is true, then it is just the beginning.

The true power of this argument lies in its extreme generality. The argument doesn't just apply to drones. It can apply to the introduction of many new technologies and for the continued use of existing ones. It shows that a major upside of military technologies (reduced casualties or reduced financial costs) comes with a major downside (increased number of wars), which may be large enough to wipe out all the benefits.

Consider an apparently benign technology, such as advances in battlefield medicine aimed at helping treat wounded soldiers. There have been many such technologies, including ambulances, amputation, and blood transfusions. These have clearly had good direct effects, saving the lives of soldiers and limiting their injuries. However, by lessening the human toll of war, they have lowered the cost of going to war and have thus probably had the effect of statistically increasing the number of wars entered by the countries deploying them.¹⁰ Presumably the effect is about the same proportionally for both drones and battlefield medicine: for each thousand lives on the deploying side that they save, this will have a similar effect on the cost of going to war, and thus a similar effect on entering more wars.

Few would even stop to entertain the possibility that battlefield medicine might be an immoral advance, but the same argument that was run against drones works equally well here. Indeed, it works even better, since my third problem for the

¹⁰ While more deaths and more injuries typically increase the political costs of a war, there is a subtlety in the relationship. It is possible that grievous or disfiguring injuries increase the political cost more than deaths as they are a more visible and lasting reminder. Thus it is at least theoretically possible that better combat medicine might raise the cost of going to war if by saving lives, it means that there are more survivors with grievous or disfiguring injuries.

argument doesn't apply: even when everyone has adopted a battlefield medicine technology, it still lowers the cost of going to war.

It is sometimes claimed that there is an important moral distinction between weapons and other means of supporting war, such as food or medicine for the combatants.¹¹ However this putative distinction regards deontological considerations about whether it is permissible to attack the people involved in the production or supply of those things. Here I'm concerned only with the morally relevant consequences of providing medicine, which will include increasing the number of future wars. These consequences will be just as real and relevant to our ethical thinking whether produced by weapons or by medicine.

One might think that I am running this argument with battlefield medicine as a *reductio* of the original argument regarding drones.¹² I am not. Modulo the concerns I raised earlier, it seems to me to be a good argument. If more people died in battle due to worse medicine, we would presumably have fewer wars. I don't know if this effect outweighs the benefits of battlefield medicine, but it does at least call its overall value into question. It also makes me worry. We tend to look for arguments like this one about changing the costs of war when the technology under consideration seems sinister right from the start. We don't look nearly so hard for such indirect effects concerning a technology that seems benign even though the consequences could be just as large. This suggests that we have major blind-spots in the evaluation of new technologies in war and could be making some serious mistakes.

The equilibrium argument is not restricted to evaluating new technologies. It can also be applied to policy changes related to war. Moreover, it can work in reverse. A case that involves both of these is that of introducing the draft.¹³ There are many negative direct effects of forcing young people into the military and it is these which are normally the main focus of debate. However, having the draft also *increases* the costs of going to war. It spreads the human cost of war across all social classes and makes it unavoidable. This lowers the voters' appetite for war, making war more politically costly. It also means that the children of the politicians are equally vulnerable, adding personal cost of declaring war to those who ultimately make that decision, in addition to the political cost. Thus introducing the draft is one way of making the country go to fewer wars. If someone were opposing the draft, they would therefore need to consider that their actions would on average lead to more war — a heavy cost indeed, and one that may not be worth paying.

We could also return to the topic of decapitation strikes and see how the equilibrium argument affects them. One consequence of such strikes — if they can be done

¹¹ (Forge 2013); (Walzer 2006, 146); though others disagree (Fabre 2009).

¹² (Strawser 2010) considers a similar argument to mine, though he claims that if the cost of going to war argument concludes that drones are overall immoral, it will also conclude that *all* other military technologies are immoral. At one point (p. 359) he suggests this may be true and at another (p. 361), that it would constitute a *reductio* of the argument against drones.

¹³ I owe this point to Andrew Snyder-Beattie.

successfully — is that the loss of life to both sides in the war is much smaller. Indeed this is putative benefit of decapitation. However, this may lower the cost of going to war, leading to more war. Alternatively, since decapitation strikes would open the country's leadership up to attack, they may raise the personal cost of going to war among these decision makers and *raise* the cost of war, leading to fewer wars. It is difficult to determine which effect would dominate.

An even more striking example is due to George Orwell. In two newspaper columns written near the end of the Second World War, he surprisingly mounted several arguments to undermine the very reasonable sounding position that we must stop bombing civilians.¹⁴ One of these was an equilibrium argument: that the immunity of civilians in war was one of the very things that had made modern war possible. He saw war as hellish and reasoned that civilians could only possibly support it if they were shielded from it. Drop that shielding and the political costs of war would rise sharply enough to greatly reduce the number or length of wars and their overall damage.

Thus the real power of the equilibrium argument is to show that when it comes to assessing the morally relevant consequences of introducing new military technologies, we might be completely at sea. The direct short term effects of the change that are so salient to us might be frequently overshadowed by unintended long run effects that frequently run in the opposite direction. This dynamic is well understood in economics, and forms the core of Frédéric Bastiat's famous essay of 1850, 'What is seen and what is not seen':

'...an act, a habit, an institution, a law produces not only one effect, but a series of effects. Of these effects, the first alone is immediate; it appears simultaneously with its cause; *it is seen*. The other effects emerge only subsequently; *they are not seen*; we are fortunate if we *foresee* them.

There is only one difference between a bad economist and a good one: the bad economist confines himself to the *visible* effect; the good economist takes into account both the effect that can be seen and those effects that must be *foreseen*.

Yet this difference is tremendous; for it almost always happens that when the immediate consequence is favorable, the later consequences are disastrous, and vice versa. Whence it follows that the bad economist pursues a small present good that will be followed by a great evil to come, while the good economist pursues a great good to come, at the risk of a small present evil.'¹⁵

I do not mean to claim here that the unseen indirect effects always dominate in the ethics of war — that they are the submerged part of the iceberg, that they always reverse the direction. This may be so, or it may be that they are typically less than the direct effects and merely dilute the apparent goodness or badness of the change.

¹⁴ (Orwell 1944a) and (Orwell 1944b).

¹⁵ (Bastiat 1850, 1).

What I am claiming is that they are of large and unknown size, that they often push in the opposite direction, and that in some cases they may swamp the direct benefits. That contrary to initial appearances, we might currently have almost no idea about the overall consequences. And that it might be a great mistake to spend so much of our time coming to agreement over that which is seen, while mostly ignoring that which is not seen.

I am not making the familiar argument that consequences can be extremely difficult to predict, so moral evaluation of them is futile.¹⁶ I think that there are good responses to that involving the balance of probabilities.¹⁷ If these shadowy indirect effects were known to be impossible to evaluate *ex ante*, we'd be right to ignore them. But, it appears that sufficient details of their overall shape and size might yield to further analysis and, if so, such analysis is urgently required.

I've focused on arguments about increasing or decreasing the cost of going to war, but it is worth noting that equilibrium arguments can go beyond this. For a familiar example, suppose a country is considering deploying poison gas in a particular battle. This may help win that battle, but is likely to also shift the equilibrium for the entire war from one where chemical weapons are never used to one where they are deployed when convenient. One could see this as lowering the opposing force's cost of using poison gas. Previously its use would have come with high political cost at home and risk of reprisals in kind from their adversaries. But after these adversaries have already used poison gas, there would be more public support at home for its use and it would no longer be clear that refraining from using gas would decrease the chance of their adversaries using it.

There are also other arguments about indirect long-term effects that go beyond equilibrium effects. For example, it is plausible that the use of drones would make it more likely for us to see unmanned aerial vehicles with no human in the loop. If the deployment of such fully autonomous robotic weapons would be sufficiently bad and the deployment of drones sufficiently likely to lead there, then that could constitute an indirect long-term argument against deploying drones. Such arguments should also draw our attention, though unlike the arguments about the cost of going to war, it is not clear that there is a systematic pull towards the indirect effect being opposite in value to the direct effect.

Conclusions

We have seen how two theoretical tools — counterfactuals and equilibria — can help us understand the moral consequences of deploying drones in war. Each of them has the power to overturn one's initial impressions of the moral value of drones. Counterfactuals show us that it doesn't matter if drones are bad: what matters is whether they are better than what they would replace. They also show us that we should ideally be assessing more fine-grained claims — 'should we ban the use of

¹⁶ (Lenman 2000).

¹⁷ (Burch-Brown 2014); (Cowen 2006).

drones for targeted killing?’ rather than ‘should we bad drones?’. The potential for introducing a new technology to shift a key equilibrium has shown us that there are often systematic indirect effects of our actions which could outweigh all or most of the direct effects.

In both cases, they suggest that the evaluation of new technologies in war is much harder than it first appears. They show that whether the consequences are overall good or bad turns on empirical questions that are very difficult to answer (such as ‘to what extent will drones crowd out bombing?’ and ‘how much more likely are we to go to war when its costs are reduced by a certain amount?’¹⁸). We cannot simply ignore such effects and such questions, as we have seen that they may have a very large effect on the evaluation, and this effect may be in a systematic direction. We often won’t have the skills to begin to address such questions, but we need to start talking to those who do: in particular, to economists. Until we fully incorporate tools such as counterfactuals and equilibria into our thinking, and until we make some progress on the underlying empirical questions, I worry that we really may be completely at sea when it comes to evaluating new military technologies.

References

- Bastiat, Frédéric. 1850. ‘What is seen and what is not seen’, in *Selected essays on political economy*, George B. de Huszar, ed. New York: D. Van Nostrand Company, 1964, pp. 1-50.
- Beauchamp, Zach & Savulescu, Julian. 2013. ‘Robot guardians: teleoperated combat vehicles in humanitarian military intervention’. In Bradley Jay Strawser (ed.) *Killing by remote control: the ethics of an unmanned military*. (Oxford: Oxford University Press), pp. 106-125.
- Brough, G. and D. Williams, A. 2011. ‘Decapitation strategy’. *The encyclopedia of war*.
- Burch-Brown, Joanna M. 2014. ‘Clues for consequentialists’, *Utilitas* 26:105–19.
- Cowen, Tyler. 2006. ‘The epistemic problem does not refute consequentialism’, *Utilitas* 18:383–99.
- Fabre, Cécile. 2009. ‘Guns, food, and liability to attack in war’, *Ethics* 120:36–63.
- Forge, John. 2013. *Designed to kill: the case against weapons research*. (Dordrecht: Springer).

¹⁸ In economics, this would be called the *elasticity* of going to war with respect to cost. This is a parameter which is typically between zero and one. The value represents the ratio between the percentage change in the dependent quantity (e.g. number of wars) and the percentage change in the independent quantity (e.g. cost of going to war). Economists have many theoretical and statistical tools for estimating such elasticities.

- Grut, Chantal *et al.* 2012. *Counting drone strike deaths*. Human Rights Clinic, Columbia Law School.
<http://web.law.columbia.edu/sites/default/files/microsites/human-rights-institute/files/COLUMBIACountingDronesFinal.pdf>
- Human Rights Watch & International Human Rights Clinic. 2012. *Losing humanity: the case against killer robots*. (New York: Human Rights Watch).
http://www.hrw.org/sites/default/files/reports/arms1112ForUpload_0_0.pdf
- Johnston, Patrick. B. 2012. 'Does decapitation work? Assessing the effectiveness of leadership targeting in counterinsurgency campaigns', *International Security*, 36:47-79.
- Jordan, Jenna. 2009. 'When heads roll: assessing the effectiveness of leadership decapitation', *Security Studies*, 18:719-755.
- Lenman, James. 2000. 'Consequentialism and cluelessness', *Philosophy and Public Affairs* 29:342-70.
- New America Foundation. 2014. 'Drone wars Pakistan: analysis'.
<http://natsec.newamerica.net/drones/pakistan/analysis>
- O'Connell, Mary Ellen. 2012. 'Unlawful killing with combat drones: A Case Study of Pakistan', in Simon Bronitt, Miriama Gani, Saskia Hufnagel (eds.) *Shooting to Kill: Socio-Legal Perspectives on the Use of Lethal Force*. (Oxford: Hart), pp. 263-292.
- Orwell, George. 1944a. 'As I please', *Tribune* (19 May 1944), Volume III 179-81, in *The Collected Essays, Journalism and Letters of George Orwell*.
- Orwell, George. 1944b. 'As I please', *Tribune* (14 July 1944), Volume III 213-5, in *The Collected Essays, Journalism and Letters of George Orwell*.
- Orwell, S & Angus, I (eds.). 1970. *The Collected Essays, Journalism and Letters of George Orwell*. (Hammondsworth: Penguin).
- Plaw, Avery. 2010. 'Sudden Justice'. 7th Annual Global Conference on War and Peace, Prague, 1 May.
- Rawls, John. 1971. *A Theory of Justice*. (Cambridge, Massachusetts: Harvard University Press).
- Russell, James. 2013. 'America's leadership decapitation policy is prescription for endless war'. *The Guardian*, February 12.
- Serle, J. & Ross, A. K. 2014. 'February 2014 Update: US covert actions in Pakistan, Yemen and Somalia', The Bureau of Investigative Journalism, March 4.
<http://www.thebureauinvestigates.com/2014/03/04/february-2014-update-us-covert-actions-in-pakistan-yemen-and-somalia/>
- Strawser, Bradley J. 2010. 'Moral predators: the duty to employ uninhabited aerial vehicles', *Journal of Military Ethics* 9:342-68.

Tzu, Sun. 1963. *The Art of War*. Samuel B. Griffith, ed. (Oxford: Oxford University Press).

Walzer, Michael. 2006. *Just and Unjust Wars*, 4th ed. (New York: Basic).

Whetham, David. 2013. 'Killer drones: the moral ups and downs', *The RUSI Journal*, 158:22-32.

Wright, Robert. 2010. 'The price of assassination', *The New York Times*, April 13.