Philos. Technol.
https://doi.org/10.1007/s13347-017-0299-6

RESEARCH ARTICLE

Just War, Cyber War, and the Concept of Violence

Christopher J. Finlay[1]

Received: 4 December 2016 / Accepted: 20 December 2017

Abstract Recent debate on the relationship between cyber threats, on the one hand, and both strategy and ethics on the other focus on the extent to which 'cyber war' is possible, both as a conceptual question and an empirical one. Whether it can is an important question for just war theorists. From this perspective, it is necessary to evaluate cyber measures both as a means of responding to threats and as a possible just cause for using armed kinetic force. In this paper, I shift the focus away from 'war' as such in order to ask whether some cyber threats might justifiably be characterized as a form of 'violence.' Some theorists argue that the term violence ought to be defined so as to encompass things like 'structural' harm or harm by neglect and thereby question implicitly the focus of just war theorists on armed force. This paper draws on a theory of violence I developed elsewhere as a defence of just war theory's narrow understanding of violence. According to the 'Double-Intent' theory, a distinctive form of 'Violent Agency' is the factor uniting the category of violence while partly accounting for the peculiar moral connotations of the term. Here, I argue that the resulting definition of violence reshapes the category in a way that includes some forms of cyber-attack. This may help us to see where cyber might fit in relation to just war theory and the ethics of kinetic attack.

Keywords Cyber-attack · Cyberwar · Violence · Deterrence · Just war theory

1 Just War and Cyber War

What guidance, if any, should the just war theory be able to offer about cyberwar? Some scholars argue that the very term 'cyberwar' is a misnomer and that identifying

This paper is based on a talk presented at a workshop on 'Landscaping Strategic Cyber-Deterrence,' Oxford Internet Institute, 3 June 2016. My thanks to Mariarosaria Taddeo for the invitation to participate and to the other participants for their engagement with the paper, particularly, Joseph Nye, Jr., Mervyn Frost, Jeff McMahan, and Paul Schulte.

✉ Christopher J. Finlay
Christopher.j.finlay@durham.ac.uk

[1] School of Government and International Affairs, Durham University, Durham, UK

cyber-attacks with war reflects basic misconceptions about both. Yet, the advent of cyber-threats of various kinds as a fact of international security has led governments, expert bodies, and scholars to incorporate these phenomena within a just war framework.[1] In the *Tallinn Manual on the International Law Applicable to Cyber Warfare*, for instance, the International Group of Experts state that 'both the *jus ad bellum* [concerning justifications for initiating war] and the *jus in bello* [concerning just conduct *in* war] apply to cyber operations' (Schmitt 2015, v). As a result, the *Manual* tracks the logic of the 1977 Additional Protocols to the 1949 Geneva Conventions, treating cyber measures as a new category of 'weapon' and therefore as being subject to the requirements and restrictions of International Humanitarian Law (henceforth, IHL; ibid., v-vi). This ties cyber security and cyber-attacks closely to two of the main subdivisions of just war theory (JWT), *jus ad bellum* and *jus in bello*.

On the face of it, this seems like it might be the right direction to take in thinking about the ethics of cyber-attacks and how to respond to them. The degree and, arguably, the type of destructiveness that some cyber-attacks are capable of inflicting suggest that the sorts of response they might justify could be similar to those warranted by 'kinetic' attacks by means of conventional armed force in a range of possible cases. If cyber-attacks can significantly deteriorate the effectiveness of either military technologies or peacetime infrastructure, for instance, or even cause physical harm to individuals, then they appear likely, on the face of things, to justify more conventional defensive or retaliatory measures in imaginable circumstances. And in a world that is increasingly reliant on and, indeed, composed of informational structures, artefacts, institutions, and resources, a sufficiently serious attack whose destructiveness was limited to damaging information as such in the cyber-domain could also, one would imagine, pass this threshold of severity too. The USA, for instance, has adopted the principle that retaliation against cyber-attacks may take the form not only of cyber-counter-attack but also attack by conventional military means (Obama 2011, 14). By the same token, if cyber-attacks are sometimes the means of deterring hostile actions of one sort or another, then they too may be subject to stringent justificatory demands akin to those governing conventional armed attack.

Trying to specify exactly what it means to place cyber-attacks within the map of just war concepts and criteria is not a simple matter, however, and raises a variety of further questions as scholars have recognized (see Taddeo (2014), 37–8). The view proposed in this paper is that, if some cyber-attacks are comparable to some kinetic attacks, and if JWT is the right ethical-legal framework for addressing normative questions about kinetic attack, then this would also seem to be the right framework within which to address at least some questions about the ethics of cyber security. But it will not be the right place to discuss *all* such questions. My argument develops the point that only some types of cyber-attack are equivalent to armed, kinetic attack, and give prima facie warrant for armed, kinetic defence. And for the same reason, these types of attack would be peculiarly hard to justify, just like the use of conventional weapons. I will call this subset of cyber threats 'Violent Cyber Attacks.' This leaves a wide range of types of cyber-attack outside the remit of JWT.

To make the case, I argue in Sect. 2 that attempts to locate the ethics of cyber security in relation to JWT are stymied by a common tendency to assimilate cyber

[1] For sceptical views, see Rid (2013) and May (2015).

threats by invoking the concept of *war*. Instead, theorists of cyber security should focus on the notion of *violence* and follow the trend of recent JWT in moving away from traditional notions of war. In Sect. 3, I examine the role that the concept of violence plays in the ethics of armed conflict and then outline in Sect. 4 an approach to its definition that encompasses some cyber-attacks. In Sect. 5, I indicate some distinctions that this analysis suggests might be important within the category of 'cyber-attacks' before indicating in Sect. 6 how these kinds of action might be incorporated within the various dimensions of JWT.

2 The Concepts of War and Violence

It might seem probable that one could integrate the concerns of cyber security into the just war theory by coupling them together using the concept of 'war': JWT is concerned with war; cyber threats are thought by some to amount to a form of war and may contribute to the waging of war in a conventional sense (Durante 2015, 369–70); and since both deal with war, JWT should be able to apply quite directly to 'cyber war.' But I doubt this can really help clarify the place that cyber should have in JWT because not only is the applicability and meaning of the word war quite controversial in relation to cyber-attacks but it is also arguably quite doubtful and unclear what it means even in JWT.

The first problem with a war approach is that theorists disagree as to whether cyber-attacks themselves can constitute a form of warfare or whether a coordinated set of such attacks could be interpreted as comprising a war taken together. Thomas Rid, for example, has argued that the fact that cyber threats do not (or do not directly) involve 'violence' in a conventional sense, by which he means using physical force to cause potentially lethal harms, is one reason to doubt it (e.g. Rid 2012; cf. Stone 2013). On this sort of view, then, JWT ought to regard cyber-attacks as 'measures short of war'—*far* short in fact—and therefore subject to a different set of guidelines. Others argue, by contrast, that the prevalence of cyber-attacks as a supplement to the kinetic forces available to states 'means that a deeper comprehension of what war is in the cyber age' is needed (Durante 2015, 370).

Further complicating the issue is the fact that the putative connection between cyber-attack and war is imagined (and sometimes experienced) in different ways, each suggesting a different relationship with war and pointing towards different potential implications for the way we might work out the ethical and legal ramifications of the phenomenon. What we might call a pure cyber war, for instance, might involve conducting an international conflict purely by means of reciprocal cyber-attack. More probable, perhaps, is a practice of war in a more conventional sense that increasingly involves cyber-attacks as a supplement to conventional weapons (Durante 2015). Cyber-attacks can be used in such a way as to affect a state's ability to wage conventional war quite directly, for instance, though this might occur during an established state of war or as a warlike act in peacetime—such as in the case of Stuxnet when it was used to attack Iranian nuclear capacity building. Moreover, there are at least hypothetical—perhaps plausible—scenarios in which the threat from a cyber measure is commensurate with some forms of armed kinetic attack, even if it is not aimed directly at war-making capabilities or launched during an armed conflict.

But even without the pressure of these new phenomena, the concept of war has in any case become more widely contested and increasingly indistinct. Up to the mid-twentieth century, the dominant modern understanding of war resembled Jean-Jacques Rousseau's: war was a formally recognized condition of enmity between states with clearly defined limits in terms of who was and who wasn't involved, in what capacity, and over what period of time or tract of geographic space (Rousseau 1762 / 2004, 10; Walzer 1977). But over the last generation or so, just war theorists have increasingly tended to think about war in a much more open-ended way: it encompasses small-scale conflict possibly running from the case of two people fighting in the street (one an aggressor, the other defending themselves) on a moral continuum that runs up to much more organized, coordinated large-scale uses of violence by lots of individuals (McMahan 2004; cf. Fabre 2012).

A decreasing emphasis on a formal conception of war has also been a feature of contemporary international law. Since World War II, jurists have addressed the problem of encompassing the right range of conflicts by making International Humanitarian Law apply to 'armed conflicts' rather than war as such. This facet of post-World-War-II international law marks a shift away from a paradigm in which states were the only—or at least the predominant and paradigmatic—parties to war and according to which wars came into being when sovereign powers declared them. But the category of armed conflict is if anything even less promising as a way of unifying a field that includes both harms inflicted by conventional military means and cyber-attacks. Apart from anything else, the term 'armed' is question begging in this context.[2]

So, I do not think war as a concept (or armed conflict) serves as a good, basic starting point for figuring out the place of cyber in relation to JWT. It neither helps show that (some or all) cyber-attacks clearly *belong* in JWT nor does it show that they necessarily *do not*. But we might get further by thinking about another concept: that of violence. What JWT and international law do seem to concern themselves with along that continuum of conflicts are the troubling questions of how to respond to threats from a spectrum of intentional, destructive harms denoted by that term and second, whether, when and how to *use* violence as a means of responding. So, if this is the common thread running through the various concerns of JWT, then we might get further in thinking about the place that cyber threats ought to have in the ethics of war by investigating the concept of violence. We would have to ask (1) what the term violence actually means or ought to be understood to mean and (2) whether some cyber-attacks might not themselves sometimes amount to a form of violence or have enough features in common with paradigm cases of violence to be directly comparable with violence.

To begin this exploration, I turn now to the concept of violence and the peculiar normative connotations of the term. These are important in explaining the way violence is given special treatment—both as cause and means—in JWT. Identifying these connotations helps specify the significance that defining some cyber-attacks as acts of violence could have for the ethics of cyber security and strategic cyber-deterrence.

[2] The appropriateness of the term 'cyber armed conflict' is suggested in relation to the applicability of IHL, for instance, by Schmitt and Vihul (2016, 35).

3 Violence

JWT's focus on the ethical problems of defending against wrongful violent threats and of using physical violence to do so is motivated by the peculiar normative significance commonly attributed to the various phenomena associated with the concept of violence.

First, as Hannah Arendt argued, it is commonly assumed that when human affairs turn violent, particularly in politics, they in some sense come to be 'set … apart from' other phenomena (Arendt 2006, 8). So violence is somehow especially troubling in a way that other forms of action or behaviour are not and a step over the threshold into the use of violence is something we mark as significant and worrying. Second, and apparently as a consequence, those who resort to violence carry a peculiarly weighty responsibility to justify themselves: as Kai Nielsen puts it: '[p]olitical violence, like violence generally, is in need of a very special justification indeed (1982, 25). Actions identified as violent in the relevant sense, on this view, are implicitly seen as either prima facie or pro tanto wrong: it may be possible to justify them, but only given certain exceptional circumstances and for especially important ends.[3] A third feature of the idea is that wrongful violence has a special role in triggering justified violence. So it is something along the lines that the paradigmatic justification for violence is *violence itself*. Or, at least, identifying something as violence has a tendency to make counter-violence appear justifiable.

As regards these features of the idea of violence, for present purposes, therefore, three questions arise: (1) Why is violence seen as having these connotations? (2) Are the reasons for this belief philosophically sound and are they sufficient to justify its special status? That is, do paradigm cases of violence, whatever that is, display features that could account for these normative connotations? And if we should take these connotations seriously about the notion of violence, then it suggests that the question of whether cyber-attacks sometimes constitute a form of violence in a literal sense is important: if so, then violent cyber-attacks are likely to have the same moral significance. So (3), can they? And, if so, might 'cyber violence' have similar ethical connotations and consequences?

The question of how violence is—and how it ought to be—defined has proven controversial, and a variety of competing views attempt to challenge conventional wisdom in different ways. The popular understanding that much debate is pitched in relation to is sometimes called the 'strict' account since it gives rise to a relatively narrow category of action types. Associated with actions such as shooting, punching, stabbing, and blowing things up, it is generally thought to exclude the range of things encompassed by 'cyber-attack.' It might be defined as follows:

> The Strict Definition of Violence: the intentional infliction of (severe) harm by human agents on others usually effecting itself in physical injury in paradigm cases but (on some accounts) also encompassing psychological damage. Acts of

[3] A relative of this view even takes violence to be wrong by definition: Robert Paul Wolff offered an influential if rather peculiar argument in 1969 premised on what he took to be a frequent tendency to differentiate between 'force' which could be 'legitimate' and violence which was defined in contradistinction to it.

violence are typically also descriptively violent in that they are sudden, forceful, and sensational. Harms are inflicted by directing such actions towards either a victim's body or something they value (such as their property).[4]

On this understanding, violence has three important features:

1. Agency: it involves inflicting harm through intentional action.
2. Bodily Harm: variants of the definition tend to treat bodily—physical—harm as paradigmatic. Other possibilities are included insofar as they appear similar or are in some sense derivatives or close relatives of bodily harm, e.g. psychological harm to the person or damage to property.
3. Violence in a Descriptive Sense: and mediating between them is the expectation that 'acts of violence' (as John Harris puts it) will also be 'violent acts' in some sense. They will be sudden, perhaps loud, and forceful.

This encapsulates, I think, a fairly common way of defining violence implicit in ordinary speech in Anglophone contexts but also followed, more or less, in some of the analytical literature.[5] But it has come under pressure at different times from a variety of revisionary efforts by philosophers, who seek to redefine the concept and usage to reflect what they take to be a more suitable range of moral and social concerns.

John Harris, for instance, argues that the strict conception does not actually define a distinctive category of things (Harris 1980; cf. Coady 2008, 30).[6] To show this, he generates dilemmas for his readers between including cases that, on the one hand, *intuitively* seem like they ought to be categorized as acts of violence but that, on the other hand, do not have the descriptive features necessary to be defined as such on the strict understanding of the word. Consequently, they put pressure on the common definition insofar as it tracks the sensational features of the 'fire and sword' idea of violence. For instance, he offers the use of slow poison as a problem case. Surely, we might think, it has to be an act of violence to put poison in someone's food that puts them into a lethal sleep. And yet, this does not appear to have the features that an analytical ethicist like Tony Coady attributes to violence: it is simply *not violent*.[7] It involves no dramatic uses of physical force; there is no suddenness, loudness, tearing, or shattering.

[4] For this synopsis, see Finlay (2017, 71). For examples, see, for instance, the strict conception disputed by Harris (1980, 15), also Madden Dempsey (2006, 310–11) and Jacquette (2013), cf. Geras (1989, 187).

[5] For common usage, see the *Oxford English Dictionary*, where violence is defined as, 'Behaviour involving physical force intended to hurt, damage, or kill someone or something.' Note that it offers a specifically legal usage as an alternative: '*Law*: The unlawful exercise of physical force or intimidation by the exhibition of such force.' The latter is what Wolff (1969) takes violence to be. For scholarly argument along these lines, see Coady (2008).

[6] See also Galtung (1969; cf. Coady 2008, 24–9) and Garver (2009), which extends the category to include 'covert institutional violence'; and Lee (1996, 330), on 'the moral continuity between the harms of social disorder and the harms of social order.'

[7] In cases where poisoning is 'slow-acting' and requires 'repeated doses,' whose 'destructive effects are gradual and cumulative,' Coady writes, 'I suspect we should not call the poisoning a violent act' (whereas using 'swift-acting' poisons like those deployed in war would be a 'fairly clear' case of a 'violent act'). The former types of case belong, he says, to 'territory lying on an uncomfortable borderline between violence and non-violence' (2008, 41).

Michael Schmitt alludes to this problem in his discussion of the notion of 'attack' in international law and its applicability to cyber operations. Additional Protocol 1 (article 49(1)) makes clear that attack is to be understood in terms of violence in something like the strict sense of the word, excluding things like propaganda and activities with economic consequences such as embargos. Attacks, in this text, are 'acts of violence against the adversary, whether in offence or in defence' (Schmitt 2012, 286). Schmitt reads violence, here, as a term indicating the release of physical forces, but he resists an overly restrictive interpretation of the letter of the law by reference to the prohibition on chemical and biological weapons. These might be said to fall outside that definition of violence, and yet, they clearly fall under the purview of Additional Protocol 1. In which case, he concludes, the intention of the legal statute should be interpreted according to the *consequences* of the action rather than whether it is strictly speaking violent in respect to the unleashing of force:

> A careful reading of Additional Protocol I's prohibitions and restrictions on attacks discloses that the concern was not so much with acts which were violent, but rather with those that have harmful consequences (or risk them), in other words, violent consequences. In great part, the treaty's object and purpose is to avoid, to the extent possible in light of military necessity, those very consequences (Schmitt 2012, 290).

Following the line of argument pursued by Harris, however, there are two further problems down that road. First, not all weapons need have 'violent consequences,' just as not all involve the violent projection of force in their execution. Poisons, for instance, need not involve any violence in their sensory characteristics to effect their consequences. But if we take violent consequences to mean something like '*sufficiently harmful* consequences,' in a broader sense, then we end up proving too much: as Harris in particular argues (but also Johan Galtung and other critics of the strict conception), harmful consequences—often truncating or otherwise disfiguring the people's lives—arise from all sorts of anthropogenic processes and structures (Harris 1980; Galtung 1969).

We are therefore still left wanting an account of how violence can be defined in such a way as to match the category it delineates to the set of things that have the moral characteristics connoted by the term. And we also lack an account of how cyber operations relate to violence and, hence, to the notion of attack operative in IHL and JWT. If we follow ordinary usage and the letter of the law in Additional Protocol 1, then cyber-attacks appear to be excluded en bloc; but if we focus on consequences without insisting on intentions or some other distinguishing feature of the act, we include too many other things potentially as equivalents to conventional, armed, kinetic attack and, hence, potentially as just cause for defensive armed force. The question is whether a conception of violence is possible that will encompass the right range of intuitive cases while also avoiding dramatically revisionary consequences such as those that Harris and others propose. If so, then it may be possible to preserve the link between the more familiar category of acts and the moral connotations that seem important to JWT. And if it is possible to assimilate the types of case that Harris, Coady, and Schmitt found problematic in light of the Strict definition—such as

poisoning and chemical weapons—then it may also be possible to assimilate some types of cyber-attack.

The theory I will now set out in Sect. 4 can achieve these aims but with some revisionary consequences that are less dramatic than those that critics of the strict conception aim at. One of these is the suggestion that seemingly intangible attacks like those involved in cyber war might qualify either as *acts of violence themselves* or as *integral parts of acts of violence*. Either way, as I will argue in Sect. 5, it shows how certain kinds of cyber-attack might be assimilated to the category of action types that concern JWT.

4 The Double-Intent Conception of Violence

According to the 'Double-Intent Conception of Violence' (Finlay 2017), the range of cases that we associate paradigmatically with violence in ordinary speech as well as in ethics corresponds to a definition as follows:

> VIOLENCE$_i$: is defined by the presence of Violent Agency consisting of the intentional infliction of [1] *destructive* harm by human agents on targets using a technique chosen with the further intention [2] of eliminating or evading the target's means of escaping it or defending against it. In paradigm cases of violence by single-minded attackers, [2] will be realized as far as is *necessary* to secure [1] or, failing that, as far as *possible* to maximize the chance of doing so (Finlay 2017, 73).

So, the first defining feature of 'Violent Agency' and, hence, Violence$_i$ is a particular kind of 'double intention'; the second is its orientation towards 'destructive' harming. Let me go through these two steps in turn in more detail.

4.1 Double Intent

Violence is defined, on this account, first by a double intention: on the one hand to inflict harm and, on the other, to narrow the window of opportunity within which its victim can respond, to whatever extent is necessary for success and possible. To illustrate, consider the way many of the *means* of violence are designed. Someone throwing a spear in an ancient war clearly intends their target to be harmed: wounded if not killed. But the choice of weapon and the technique in which it is used realize a second, supporting intention too: this is to reduce the alternatives available to the target that might permit them to evade the harm. Whereas the sharpness of the point might be seen as a *force* multiplier, concentrating the thrust from the thrower's arm on a minute point of contact, the speed with which the spear flies serves both to communicate this force *and* as a *dominance* multiplier. By dominance multiplier, I mean that it increases the vulnerability of its target to the harm. This, it achieves by reducing the window of opportunity within which to react, whether by stepping to one side, or by parrying the blow with a shield or a sword.

Throwing a spear works in a similar way to firing a bullet from a gun. If the shooter takes aim from a vantage point and, moreover, does so in a way that is

concealed from the target, it thereby increases the degree of domination: on the one hand, the attacker's ability to inflict harm at will rises while, on the other, the target's ability to evade or parry it declines. Paradigmatic acts of violence, I argue, always combine these two intents, commonly aiming at a transient, momentary relationship of intensified dominance within which to execute the intended harm. The familiar category of action types sometimes referred to as 'kinetic violence' or 'kinetic armed force' commonly employs the factors of speed, distance, secrecy, and surprise to narrow the response window and inflict harm without impediment. The more effectively a technology does so, the better suited it is to employment as a means of violence. The combination of a chosen technology (bomb, gun, knife, or fist) with a particular method of employment (booby trapping, sniping, stabbing, or punching) may be referred to as a whole as a *technique* of violence in general and as an *act* of violence in each particular case.

Whereas, classic violent means such as shooting or punching create a radically intensified relationship of dominance in a sudden, transient way, the Double-Intent Account also encompasses the possibility of opportunistic violence. Finding someone to be in an established position of vulnerability—due to class or gender relations, for instance, or through a disabling health condition—and exploiting this fact in order to inflict destructive harm are other ways in which double intent might be realized. It is also possible to prepare a victim in advance, realizing intent (2) prior to executing number (1): rendition or kidnapping, for instance, establish the relationship of dominance necessary to render someone maximally vulnerable to torture or whatever harms the perpetrator envisages. In cases of either kind, I emphasize that agents will harness or create an asymmetry of the relevant kind as far as *necessary* for the execution of the harm or, if that is not available to them, as far as *possible* for doing so. It may be, however, that the agent will not do enough to succeed, in which case, they will commit an unsuccessful act of violence—but it will still be an act of violence, nevertheless.

The usefulness of highlighting these dimensions of the act of violence is seen when we turn to some of the cases that Harris and Coady found problematic—those that sit uncomfortably on the margins of the violent and the non-violent. Poisoning, on the Double-Intent Account, appears not as a liminal case, but as a paradigmatic example of violence. Take as a hypothetical example, the following:

> SECRET ENEMY: Susan and Mary appear to all who observe their outward behaviour to be friends. Indeed, Mary shares their conviction that Susan is a loyal companion who only has Mary's best interests at heart. In fact, Susan has long held a grudge against Mary for some past, perceived insult, and decides to murder her by slipping an undetectable poison into her tea. Once she has done so, she passes Mary the tea and chats with her until it has been sipped down to the dregs. Only once it is too late for Mary to do anything about it, does Susan tell her the truth.

If violence is defined by the presence of a double intent, as I have suggested, then this is an exemplary act of violence. Not only does Susan intend that Mary be harmed, indeed killed; but also the technique she uses to bring the harm about combines a technology that can execute it with a method that renders Mary entirely vulnerable to it. By keeping the danger secret until the poison has already been ingested, Susan deprives Mary of any means of escape or defence.

If this seemingly non-standard case that has caused such trouble for proponents of a narrow or strict definition of violence can be assimilated, then so, I argue, can others such as the use of weaponized poisons and the use of sieges. Sieges, for instance, not only use violence when they kill those attempting to flee the besieged city but also constitute acts of violence writ large. This is because encircling and threatening the citizens render exposure to starvation, dehydration, and disease inescapable. By contrast, embargos and sanctions aim at a variety of different things that are less likely to fall under the definition of violence that I offer. Arms embargos, for instance, or embargos on other materials useful for internal repression or external aggression are put in place to diminish the ability of a state to cause further harm or, perhaps more often, to prevent its ability from continuing or increasing. In common with sanctions, they might be interpreted as refusing further benefits of different kinds rather than depriving their target of present goods in such a way as to cause anything that could be interpreted as 'destructive harm' as I define it below in 4.2.[8]

There is not space to defend this part of the definition fully and I have responded to a variety of possible objections elsewhere (Finlay 2017). But I will mention one, which is the worry that defining violence in this way appears to render death and injury caused as a collateral effect of warfare non-violent insofar as it is, by definition, *unintended*. Two points are worth suggesting by way of a brief response. The first is that this is true of *any* definition of violence that specifies *intended* harming and not just of a definition that specifies a double-intent. I take intention to be part of the strict understanding and something that is widely associated with violence. Secondly, defining violence by the intentions of the agent does not preclude a description of side-effect victims as having *suffered as a result of an act of violence*. Nor does it rule out saying that the agents committed an act of violence that harmed the collateral victims. It only precludes saying that the agents *committed an act of violence against them*. So, I do not think it leads to any semantic consequences that we could not quite easily live with.

Finally, by following the strict conception in the way it highlights *intention*, I set aside the notion of 'structural violence' (Galtung 1969). I do not doubt that harms might be suffered as a result of structural factors. And neither do I deny that structures do so when they place individuals in a position of acute vulnerability and thereby render them susceptible to severe harms of one sort or another. But I see the move of characterizing as violence the suffering that structures cause as a rhetorical one. It harnesses the forcefulness that the term violence derives from its paradigm cases in which agents cause harm intentionally by certain means. As such, structural violence is essentially a metaphor and its effectiveness in highlighting unjust harms is best supported, in my view, by clarifying the paradigmatic cases and the concept lying behind them.

4.2 Destructive Harming

Some might also think it an objection to a definition that relies solely on the double intent criterion that it would include things that we do not necessarily see as forms of violence. For instance, hacking into someone's bank account and stealing funds seems like an exemplary case of violence on that basis (i.e. by exploiting or creating vulnerabilities as part of the technique for harming)—but would it be violent in the

[8] Thanks to Joseph Nye for pressing me on this distinction.

usual understanding of that term? Likewise, in a more material context, burglary might often (perhaps always) be an act of violence insofar as it creates or exploits vulnerabilities (regardless of whether weapons or threats against persons or destruction of property were involved) and then executes a robbery that harms the owner of the house. Both types of action can involve [1] intent to harm by means [2] intended to deprive a victim of the means to respond and defend. So the question is whether we should admit a much wider range of behaviours into the category of violence than is usually intended by those using the word or whether we can find a legitimate way to refine the definition that excludes these cases. This is where the differentiation between what I call 'destructive harming' and what I call 'appropriative harm' is important.

Of course, burglary and assault *do* have similarities: it might be possible to bite the bullet, so to speak, and grant that violence encompasses both—it would not be *very* counter-intuitive. And for purposes of thinking about cyber, this could conceivably be helpful too: hacking as a means of theft might be treated similarly to assault and to the use of other cyber 'weapons' to cause harm without theft. But I think the common-sense tendency to distinguish between categories here has some deeper grounding. The gist of my attempt to disentangle the two is as follows (in Finlay (2017, 77–82)). I distinguish the destructive harming characteristic of violence from 'appropriative harming,' which is typical of things like theft. On the one hand, appropriative harming occurs where the harm an assailant causes to her victim is *the same as* the benefit enjoyed as a result by the assailant—i.e. they are commensurable in kind and commensurate in scale. By contrast, where the benefit to an assailant is *very different* (incommensurable perhaps or just very different in scale), then we are more likely to interpret it as a 'violent attack' rather than some sort of theft.

So, for instance, if I broke the lock of your car, hot-wired the engine, and drove off, later selling it for cash or perhaps keeping it for use as my own car, then the harm and the benefit are of the same kind and scale: a car or the value of a car. But if instead I took a baseball bat and smashed in your car windows before setting the vehicle on fire as an act of revenge, they are not: you suffer the loss of a valuable possession (as well as the anguish or fear my actions cause); I enjoy whatever satisfaction comes from avenging a perceived wrong. So, while either would constitute a setback to your interests of the sort characteristic of 'harm' in general, it seems possible to identify something here that differentiates between two subcategories: things that operate like theft (appropriative harming) and things that operate like paradigm cases of violence (destructive harming). This does not, of course, itself suggest that violence per se is necessarily worse than theft—it all depends on other factors on this line of comparison.[9] But some further clarity on this can be achieved by looking at how the Double-Intent Account of violence explains the moral connotations of the term.

4.3 The Moral Connotations of 'Violence'

So, why and in what sense is violence, as I define it, morally troubling in a distinctive way? Why does it have the sorts of normative connotations I mentioned before and

[9] Violence is especially troubling, on my account, due to its greater propensity to guarantee harms by creating simultaneous vulnerabilities. So it might be that theft in some forms and violence in all forms share this worrying characteristic. In which case, violence is not necessarily worse than theft ceteris paribus.

how might this affect cyber-attacks if some of them turn out to belong within the category of violence?

The fearfulness of violence arises first of all from the fact that, on the Double-Intent definition, whereas other kinds of agency (and structures) have only a contingent (if, perhaps, frequent) relationship with harm, harm is essential to acts of violence and their purposeful orientation towards harming is definitional. By contrast, while structures and actions that are not specifically aimed at destructive harm in the way acts of violence are might sometimes cause comparable harms, if they do, it is incidental to their purposes, or even accidental. The destructive aspect of violence on the Double-Intent Account reinforces its first feature. Whereas, if the gain to the attacker is the same as the victim's loss, there is at least an in principle possibility of restoring the lost value to its owner, destructive harm is usually marked by the elimination of some part of the victim's sources of well-being.[10] Although violence against property might be remediable in some instances, attempts at restoring well-being can at best compensate for harm in some other way in many cases. And in the most severe cases, where violence is directed against minds and bodies, compensation is impossible: limbs cannot be restored and there can be no compensation to someone who loses their life to violence.

These characteristics in turn clarify, second, why it is peculiarly hard to justify employing violence and accounts for the need for a special ethics. Whether aimed at destructive harm to the body or at other things, Violent Agency will itself increase the probability that the harm will occur to the fullest possible extent. But more than this, we can also say that even if we hold the actual harms resulting to be equal between two different acts, one by a Violent Agent and one by an agent of another sort (reckless or negligent, for instance), then the former will usually bear a greater degree of moral responsibility and, in case of wrongful harming, be more culpable than the latter. Judgements about the presence of Violent Agency thus issue in and not from judgments about culpability. And this helps account for the liability of Violent Agents to remedial harm when their actions render defensive or retaliatory measures necessary.

The significance of this facet of violence on the Double-Intent Account is thrown into sharper relief if we contrast it with the way good or evil outcomes might be generated through what Luciano Floridi calls 'distributed agency' (DA) (Floridi 2013). DA occurs when a morally significant outcome (whether positive or negative in value) occurs as a result of individual actions in their *combined* results but without any of the individual actors committing acts that are *individually* significant. So their actions might be morally neutral due to their relatively trivial scale, falling below a threshold of moral significance. Moreover, the individual actors may or may not be conscious that their acts might contribute to a greater good or evil due to DA. Indeed, since the actors, on Floridi's account, can include artificial intelligence, they may lack the capacity for such intentions or understanding, and yet they may comprise necessary causal contributors to the moral significance of the collectively generated outcomes. As Floridi makes clear, this sort of moral analysis 'evaluates actions not from a sender but rather from a receiver perspective: actions (including MAS' [those of 'multi-agent systems'], artificial and supra-agents') are assessed on the basis of their impact on the

[10] On the moral complexities of restitution in a *post bellum* context, see Fabre (2016) chapter 5 and on the difference between restitution of goods taken and reparations for goods destroyed, see p. 117.

well-being of the environment at large and its inhabitants specifically' (2013, 732). And from that perspective, it might very well be the case that the moral significance of negative outcomes is often comparable to those of intentional—indeed, doubly intentional—outcomes arising from violence.

So, as far as impacts are concerned and the moral urgency of addressing and mitigating them, it is likely that the significance of violent harms and those arising from DA will be directly comparable. However, when it comes to the question of using harmful means, including armed force, to deflect or defend against such outcomes, it is clear that those contributing to Violent Agency are likely to be liable to levels of harm directly commensurable with those that are defended against. By contrast, those who contribute in the individually negligible ways to DA envisaged in Floridi's theory are likely in many cases to be liable to little or no defensive or preventive harm. This is because, whereas the Violent Agent is, by hypothesis, responsible for intending the evil outcome, as a result of both a high degree of causal efficacy and an especially forceful kind of intention, those who contribute to DA are responsible only for individual acts that are, in themselves, morally 'neutral.'

Finally, the idea of Violent Agency therefore also helps explain why successfully persuading people to describe a particular act as unjustified violence can potentially have a permissive effect and why there is such a close association in just war theory between just cause and prior or threatened 'armed kinetic attack' (Fabre 2012, 108–10).[11] This is because actions involving Violent Agency are those which, by definition, are intended to exclude means of evasion or resistance. The more successful the act of violence, therefore, the narrower the range of options it leaves its target. If violent acts thereby eliminate the chance to block, negotiate, or escape, then it is likely that they often leave only violence itself as the victim's remaining alternative. Violence is therefore likely in many cases to generate the conditions of necessity and proportionality that justify violent defence, including armed, kinetic violence (cf. Coady 2008, 42).

5 Cyber-Attacks and the Category of Violence

On the Double-Intent Account, some types of cyber-attack—but not all—will be defined as acts of violence.

To illustrate how the two elements of the definition—double intent and destructive harming—occur in some cyber-attacks, we only have to turn to Stuxnet, the most famous recent case. First discovered in 2010, the malware that programmers used to damage the Iranian nuclear processing plant combined two indispensable components. The first executed the intended harm: it adjusted settings in the processing plant to alter the speed with which its cylinders rotated, thereby causing serious physical damage and costing the programme years in lost progress. This clearly satisfies the destructive criterion. On the face of things, the intended outcome appears to have been a gain in security for the USA and Israel against the future development of nuclear weapons by

[11] If 'institutional injustices' come to be redefined as 'forms of violence,' Steven Lee writes, 'this would be relevant to determining whether a violent response on the part of those who are being treated unjustly is justifiable, as, under common moral notions, the violence of aggression can sometimes justify the violence of defense' (1996, 68; also, van der Linden 2012).

Iran; the harm suffered by the Iranians consisted in physical damage to their centrifuges and to the programme for developing nuclear capability, whatever its actual purposes may have been. But crucially, Stuxnet had a second component: the software used a shielding mechanism 'to circumvent and compromise digital safety systems,' making it impossible for any computer it infected to detect its presence. This was needed because otherwise the computer's security software would identify it as a virus and attempt to extract it before it did any harm. The shielding component, therefore, ensured that any computer hosting the bug remained maximally vulnerable to the harm that Stuxnet was intended to execute (Rid 2013, 43–6). Like Susan, the secret poisoner, those who launched Stuxnet therefore sought to realize a double intent: *to execute a destructive harm* by means of a technique that would also *maximize the target's vulnerabilities to it* by diminishing its opportunities for evasion or defence.

Thomas Rid has argued that cyber-attacks cannot be *acts of war* since, among other things, they generally lack the direct relationship with 'lethality' that he regard as definitional of warlike violence; moreover, '[i]n an act of cyber war,' he says, 'the actual use of force is likely to be a far more complex and mediated sequence of causes and consequences that ultimately result in violence and casualties' (Rid 2012, 9). On my theory, by contrast, some such actions can exhibit the necessary features of Violent Agency and may be interpreted, therefore, as not only having causal connections with *other* acts of violence, but also as themselves *being* acts of violence. Take, for instance, the hypothetical 'logic-bombs' that Rid imagines causing train crashes, electricity blackouts, and the collapse of air traffic systems; if they were possible, such actions would clearly aim at *destructive harm* (just as Stuxnet did) and it is presumably going to be part of their operation to evade or incapacitate the firewalls and other defences set up to resist them (Rid 2012, 9). It is therefore possible to classify at least some imaginable cyber-attacks not only as precipitators of violence but as *acts of violence themselves* on the Agency Account.

Whereas, both Stuxnet and Rid's logic bombs harness informational means to bring about physical destruction of a more immediately material kind, it is also possible for cyber-attacks with a purely informational target to satisfy the necessary conditions of violence. By contrast with the hacker who steals funds or secrets, inflicting harm appropriatively (analogously with a pickpocket or eavesdropper), one who uses computer viruses to corrupt data or delete it, making it unusable, thereby inflicts destructive harms. If she does so under cover of a shielding device making her attack undetectable until it is too late (or in any other way that makes it impossible to defend against the attack), then the attack has the same agential complexion as an act of conventional, physical violence. Whether it is commensurable with one depends, then, not on the type of action per se, but on the moral significance of what each act of violence harms: a cyber-attack that destroyed data of greater value to a larger number of human beings would be worse than an attack by physical means on property of lower value to fewer.

So, the first conclusion is that some cyber-attacks not only result in or resemble acts of kinetic violence but are actually themselves properly described as acts of (non-kinetic) violence. This challenges quite a widely accepted view about cyber-attacks (for instance, Blank 2015, 94). By highlighting the ways in which a wider category of violent acts—some kinetic, some not, but all aiming at destructive harm—is unified by a particular kind of agency based on double-intent, it is possible to see how some forms of cyber-attack but not others are violent. Attacks such as the Stuxnet virus are

designed, as I have argued, to achieve their results in the same way as secret poisons or a bullet fired from a gun: they execute an intended destructive harm while simultaneously realizing the second design of depriving their targets of an opportunity to evade it. They involve a technique that harnesses components which, in other words, aim at both harm and vulnerability to harm.

But my analysis of Violence into *two* parts suggests a second potentially important way in which cyber-attack may be violent, which is that even when they are not used to execute the first intention that defines violence, cyber-attacks may be especially useful in achieving the *second*. They seem well suited, that is, as means to eliminate defences, and might thereby occur as part of a larger 'act of violence' when accompanying more traditional, kinetic methods of harm. R. R. Dipert, for instance, suggests the following hypothetical case:

> a massive cyberattack on defences a nation has against physical attack (such as radar, spy satellites, command and control systems), would risk giving the attacked nation reason to believe that a conventional attack was imminent, and then possibly trigger the conditions of justified preemptive war…(2010, 401)

I would interpret the hypothetical rather differently. Such actions might very well foreshadow a kinetic attack, but they should not necessarily be seen only as acts preparatory to a discrete act or wave of violence: they are themselves *part of* those acts. The same may be said of the historical case in which the Israeli Air Force sabotaged Syrian radar as a means of eliminating its defences before bombing the Deir ez-Zor nuclear reactor in September 2007 (Rid 2013, 11). If the method inflicted damage on the radar site, then it might have constituted an act of violence in its own right, on my account. But even if it only disabled it temporarily, it was as much a part of the greater act of violence to which it contributed—the one directed against the nuclear base—as the dropping of the bombs was. Violence, on my analysis, is defined both by the way it eliminates (or evades) defence and by the way it inflicts harm. So even if cyber measures were less effective than gunfire and rocket attack at executing harm, they are highly adaptable to the purposes of creating conditions in which it will be impossible to evade, deflect, or deter the execution of destructive harm.

Finally, the use of ransomware might sometimes constitute a third type of violence by means of cyber-attack, at least in some instances. Ransomware that follows the pattern of the recent WannaCry attack by encrypting someone's data and then demanding that payment be made in order to unlock it might be interpreted by analogy with armed robbery. In both cases, the robber extracts payment through extortion by threatening to inflict destructive harm on something valued by the victim. The armed robber points a gun, rendering the victim maximally vulnerable to physical injury; the cyber extortionist likewise threatens to render the victim's data useless to them. And both use this *threatened violence* to inflict an appropriative harm.

I propose 'Violent Cyber-Attack' as a term to designate that subcategory of cyber-attacks that are also violent in the senses I have sketched out. But just as some types of cyber-attack, then, may be interpreted as acts of violence—or as part of an act of violence—others may be excluded from the category. Stealing secrets, whether from industry, a state, or a private person, may violate rights but is not an act of violence, whether the perpetrators hack into computers or sneak documents out of an office or

someone's house. Likewise, merely obstructing someone—whether by blocking a road or denying them service by bombarding a computer with data requests—is likely to be something other than violence, even when carried out maliciously. On the other hand, if you blocked someone from escaping a collapsing building, it arguably should be seen as an act of violence and it would be defined as such on the definition I offer. And if a distributed denial-of-service attack against various different institutions and organizations, such as occurred in Estonia in 2007, had the effect of damaging those assets themselves, caused harm to users, *or* rendered the state vulnerable to harms executed by other means (such as by kinetic attack), then it should be seen as an act of violence or part of one. If the imposition of profound vulnerability is carried out for the sake of making it clear to Estonia that it may be subject to kinetic attack at any time and should feel no sense of security against such a possibility, then it ought to be interpreted as a *threat* of violence with all the coercive potential that such threats have when their credibility is proven by a substantial downpayment.

6 Cyber-attacks and the Ethics of Force

Of course, naming a subset of cyber-attacks violent need not prejudice debate about whether or not they are *justified*, legally or morally. JWT presumes that acts of violence, however the term is understood, may be justifiable if certain conditions are met. I turn now to JWT itself and the question of where Violent Cyber-Attacks ought to be located within it.

6.1 Just War Questions

To establish the appropriate place for cyber-attacks within just war theory, we need to be in a position to address three key questions.

First, can cyber-attacks or cyber-threats justify kinetic war (or warlike kinetic acts short of full-scale war)? This question is directly triggered by the question of deterring cyber-attack.[12] If there are imaginable cases in which (a) a foreign power or organization threatens a state with wrongful cyber-attack and (b) the only means of deterring or defending against the attack is kinetic, then (c) can a cross-domain defensive or retaliatory measure be justified? If so, then it seems necessary to find a way to show how a cyber-attack could meet a requirement implicit both in JWT in particular and 'common moral notions,' which is that a resort to violence requires a special justification (Lee 1996, 68). Perhaps the criterion in JWT that most directly implies this is 'last resort': as it is usually understood, this principle reflects the intuition that violence in general, war in particular, ought to be resorted to only once all other, non-violent measures have been exhausted or are manifestly futile. But it is also implicit in the generic assumption that war must have a just cause: just cause implies not only that a wrong be threatened (or have occurred) but also that it is of the right kind. Not every wrong is sufficiently egregious in nature to warrant a resort to armed force. So there is a question here of

[12] On possible use of kinetic retaliation as deterrence, see Iasiello (2014). Michael Schmitt and Liis Vihul remark on the controversial question of 'when cyberoperations alone qualify as hostilities for the purpose of initiating an armed conflict' (in 2016, 36).

commensurability: can some cyber-attacks be conducted in a manner or on a scale that is commensurate with kinetic force?

But a second set of questions must also arise, I think, if we are coming at it from a JWT angle: can cyberwar itself be justified as a form of war? And relatedly, if it can, then how might the terms of the *jus ad bellum* (JAB) apply? Are they likely to be more permissive of cyber-attacks than of kinetic measures? Should they be? These questions along with the first are likely to have practical relevance more particularly in cases where conflict involves a *cross-domain* dimension, those where kinetic force is augmented by cyber-attacks and in which, therefore, cyber becomes an important element in increasingly hybrid forms of war. And this leads to a third question: How might the occurrence of specifically *cyber* components in *armed*, *conventional conflict* be reflected in JWT?

So, for instance, as regards the JAB, can the deflection or prevention of cyber-attack or similar form *part of* a wider just cause? And, if so, might it contribute to a case for the 'proportionality' of kinetic war? Proportionality requires that decision-makers weigh relevant benefits of the successful prosecution of war against relevant costs. Both are widely seen in JWT as restricted to certain kinds of value: gains certainly include the protection of (innocent) lives from wrongful threats and probably include things like national security in a wider sense against territorial incursion and violations of sovereignty; costs are usually calculated in terms of harms to at least the most basic human interests (paradigmatically life and bodily integrity). If defence against cyber threats is seen as part of the case for claiming 'just cause' for war, then it would imply that diminishing such threats is a relevant gain and can offset relevant costs such as the loss of innocent lives through collateral harms. So this would place cyber security, as it were, on the 'plus' side of the balance sheet that the proportionality requirement demands we review before resorting to war. But it might also then have a place too on the 'minus' side. If the relevant gains are measured partly in terms of innocent human lives or the defence of national sovereignty against international aggression, then to what extent could the sort of damage that cyber-attacks launched by one's enemies in such a war count against them? At what point might the costs of such measures in our war rise to such a level as to render it disproportionate, in spite of the fact that it aimed at securing human lives from lethal, kinetic threats or over sovereign independence from the aggression of foreign powers? In the absence of a common denominator between (some) cyber-attacks and armed attacks by conventional means, it is not clear how such trade-offs would be made—or, indeed, whether it is even possible to imagine ways to think them through.

Similarly, problems then arise of how cyber measures should be compared with kinetic armed attack in the *jus in bello* (JIB). A kinetic or hybrid war begins, let us imagine, and your side has just cause. But your political and military commanders now have further decisions to make, some involving choices between strategies or tactics based on cyber and those involving conventional weapons. Presented with such decisions, should they always prefer cyber-attacks to kinetic alternatives where both have a reasonable (or possibly an equal) chance of success? Or might kinetic measures be equally justifiable or sometimes easier to justify?[13] As in the JAB, being able to

[13] Again, see Schmitt and Vihul (2016, 36) See also Michael Schmitt (2012).

make ethically defensible decisions *in bello* requires knowing whether the different kinds of measure may be compared directly and in what circumstances one might be chosen in preference to the other.

6.2 Just War Answers

If we assume that cyber-attacks are not acts of violence, then answering these questions seems relatively straightforward: no, you cannot offset anticipated side-effect deaths with reference to securing cyber assets; no, you need not count cyber damage as a cost in estimating the proportionality of pursuing legitimate war aims; and yes, uncontroversially, it is always better to damage a computer than to injure or kill a person, so cyber measures should always be preferred to kinetic in the framework of the *jus in bello*. But if I am right to argue that some cyber-attacks actually *are* acts of violence or essential components of such acts, then it suggests these questions need to be answered in a different way, one that is more sensitive to the specificities of different types of attack, whether cyber or kinetic.

If Violent Cyber-Attacks have the features of violence, as I define it, then they will also have the normative features associated with the concept of violence. Therefore, all else being equal, they can be treated in ethics the same way other cases of violence are—including kinetic violence. Of course, it might very well remain the case that the kind or quantity of harm that cyber-attacks cause as compared with kinetic arms will typically be different: that is, the set of harms that cyber-attack can cause is likely to overlap with the set of those that kinetic violence can but not entirely; they will not be wholly congruent. But in the overlapping cases, like can be treated as like. In which case, it seems likely that kinetic, armed force could be used to defend against cyber threats—and vice versa.

Whether or not kinetic defensive measures may be used against a Violent Cyber-Attack in any particular case cannot therefore be decided in advance on the basis of the way the two types of threat are characterized—kinetic or cyber—but must be judged on the basis of three factors: first, the degree and type of harm that each threatens (i.e. questions of proportionality); second, the comparison between different alternatives as to the balance of costs and benefits they are likely to bring about (necessity); and third—crucially—discrimination. So, as far as justification is concerned, the devil is in the details (and not in broad strokes description). The first two criteria may be applied simply by considering the sorts of destructive harm that will arise from the different actions being compared: let us say, between cyber-attacks that strategists in the defence ministry of a state wish to defend against and deter and kinetic measures that might be effective in responding to them and achieving those goals. If such kinetic measures inflict harms that are not disproportionate to those threatened by cyber-attack, then they can satisfy proportionality. And if the ratio between harms caused in kinetic defence measures and cyber-attacks deterred or defended against is optimal when compared with other proportionate measure that might be considered as alternative modes of response, then those measures may be said to satisfy conditions of necessity. The remaining question, then, is whether a kinetic response to cyber-attacks is likely to satisfy a criterion of discrimination.

This question needs to be addressed both in principle and as a practical matter. In principle, the account I give of how some cyber-attacks amount (or contribute) to acts of violence shows how it is possible that proportionate kinetic defence that inflicts no more relevant harms than are necessary to diminish threats from cyber-attackers *can* also meet a fairly demanding discrimination requirement. Discrimination, I shall assume, demands (1) that necessary harms be intentionally directed *only* at those who are liable to them; (2) liability is a function of moral responsibility for the wrongful threats to which defensive measures respond.[14] Violent Cyber-Attack is defined, as I have argued, by the double-intentions of an agent. Those responsible for such actions, it can therefore be argued, are morally responsible for their outcomes in a uniquely intense way: these outcomes were not merely foreseeable and nor are they intended in a simple sense; the outcomes, we might say, are 'hyper-intentional effects.' That is, they are effects that are not only intended as outcomes but guaranteed as far as possible by engaging in techniques that also eliminate key factors that could possibly impede them. Provided the perpetrators of such attacks may be identified and that the defensive measures chosen can be targeted in such a way as to discriminate between those responsible and bystanders in an appropriate way, it appears that they may satisfy the demands of discrimination in principle. Information about the agents of cyber-attacks, however, is peculiarly difficult to source and states wishing to respond in a manner that might deter future attacks are likely to be thwarted in many cases by the 'attribution problem.' So whether kinetic measures may be used defensively in a manner that satisfies the JW requirement of discrimination *in practice* may be quite a different matter.

Finally, whether cyber-attacks may be justifiable as defensive measures against either other cyber-attacks or kinetic attacks will depend on a similar ethical analysis. If the cyber measures under consideration are also Violent Cyber-Attacks, then the destructive harm they are intended to execute along with the vulnerabilities they exploit or open up as means of ensuring their success ought to be both proportionate to the threats of harm they defend against and necessary. And they must also be discriminate in distributing their intended harms in such a way as to respect the immunity of the innocent. Likewise, where Violent Cyber-Attack occurs *within* a wider cross-domain conflict or war that combines both cyber-attack and kinetic warfare, the JIB criteria will apply. *In bello* proportionality will demand that whatever measures are used, they should not cause collateral harms to civilians that are excessive by comparison with the military advantage they are expected to secure. And all such measures should satisfy discrimination so far as intended targets are concerned. Satisfying these conditions will not, on the account offered here, mean that Violent Cyber-Attack ought always to be preferred to kinetic violence: it is possible in principle that kinetic measures will be preferable in some instances, whether on the grounds that the harms they inflict are less severe or that they are targeted more discriminately and with lower rates of collateral damage.

[14] In this regard, I follow the analysis in general terms defended by Jeff McMahan. See, for instance, McMahan (2009).

7 Conclusion

The question with which I began concerned whether and how cyber-attacks or 'cyber war' ought to be located in relation to JWT. How we answer it is likely to have significant practical ramifications with the potential to affect directly the policies justifiable for governments and to have wide-ranging effects on matters such as the status of those accused of engaging in hostile cyber-operations during a conflict (see, for instance, Blank (2015, 101)). The key problem in incorporating cyber-attacks within JWT is to find a common denominator between them and the sort of armed, kinetic attack that the theory traditionally deals with. I argued that this cannot be war, not principally because of the objections that some have raised to the idea of cyberwar itself, but due to the fact that JWT presently concerns itself with a wider range of different types of conflict, some a great deal less formalized than war in a traditional sense. More promising, I have argued, is the concept of violence. To make good this claim, I offered the 'Double-Intent' account of the definition of violence. This, I argued, not only captures what is distinctive about the range of action types that are classically associated with the term violence—things like shooting a gun or throwing a punch—but also identifies something distinctive about a subset of cyber-attacks. These are attacks which combine the intention to inflict destructive harm (whether on information or on physical bodies and structures) with the intent to maximize the victim's vulnerability to such harm. Cyber-technologies can achieve this either on their own, after the fashion of Stuxnet, or in combination with other, non-cyber-technologies, as in the case of the 2007 Israeli Air Force attack in Syria. In both cases, cyber-attacks correspond closely to paradigmatic practices of violent attack and, as such, are commensurable with such attacks. Whether armed force may justifiably be used in defence against such cyber-attacks in particular contexts may be settled, therefore, on the basis of comparing the degrees and types of harm threatened by the attackers and the defenders respectively, and the extent to which defensive measures are discriminate.

References

Arendt, H. (2006). *On revolution*. New York: Penguin.

Blank, Laurie R (2015) Cyberwar versus cyber attack: the role of rhetoric in the application of law to activities in cyberspace, in Jens DavidOhlin, KevinGovern, and ClaireFinkelstein (eds) *Cyberwar: law and ethics for virtual conflicts,* Oxford: Oxford University Press.

Coady, C. A. J. (2008). *Morality and political violence*. New York: Cambridge University Press.

Dipert, R. R. (2010). The ethics of cyberwarfare. *Journal of Military Ethics, 9*(4), 384–410.

Durante, M. (2015). Violence, just cyber war and information. *Philosophy and Technology, 28*, 369–385.

Fabre, C. (2012). *Cosmopolitan War*. Oxford: Oxford University Press.

Fabre, C. (2016). *Cosmopolitan peace*. Oxford: Oxford University Press.

Finlay, C. (2017). The concept of violence in international theory: a Double-Intent Account. *International Theory, 9*(1), 67–100.

Floridi, L. (2013). Distributed morality in an information society. *Science and Engineering Ethics, 19*, 727–743.

Galtung, J. (1969). Violence, peace and peace research. *Journal of Peace Research, 6*(3), 167–191.
Garver, N (2009) What violence is. In: Bufacchi, V (ed.) (2009)*Violence: a philosophical anthology.* Basingstoke: Palgrave Macmillan.
Geras, N. (1989). Our morals: the ethics of revolution. *The Socialist Register, 25*, 185–211.
Harris, J. (1980). *Violence and responsibility*. London: Routledge & Kegan Paul.
Iasiello, E. (2014). Is cyber deterrence an illusory course of action? *Journal of Strategic Security, 7*, 1.
Jacquette, D. (2013). Violence as intentionally inflicting forceful harm. *Revue Internationale de Philosophie, 67*, 293–322.
Lee, S (1996) 'Poverty and Violence', *Social Theory and Practice, 22*(1): 67–82.
Madden Dempsey, M. (2006). What counts as domestic violence? A conceptual analysis. *William and Mary Journal of Women and the Law, 12*(2), 301–333.
May, L. (2015). The nature of war and the idea of "cyberwar". In J. D. Ohlin, K. Govern, & C. Finkelstein (Eds.), *Cyberwar: law and ethics for virtual conflicts*. Oxford: Oxford University Press.
McMahan, J. (2004). War as self-defence. *Ethics and International Affairs, 18*(1), 75–80.
McMahan. (2009). *Killing in war*. Oxford: Clarendon Press.
Nielsen, K. (1982). Political violence and ideological mystification. *Journal of Social Philosophy, 13*(2), 25–33.
Obama, B (2011) International Strategy for Cyberspace: Prosperity, Security and Openness in a Networked World, available at: https://obamawhitehouse.archives.gov/sites/default/files/rss_viewer/international_strategy_for_cyberspace.pdf (accessed, 15 April 2017).
Rid, T. (2012). Cyber war will not take place. *Journal of Strategic Studies, 35*(1), 5–32.
Rid, T. (2013). *Cyberwar will not take place*. London: Hurst.
Rousseau, J-J (1762 / 2004) *The Social Contract,* tr. Maurice Cranston. London: Penguin.
Schmitt, M. (2012). "Attack" as a term of art in international law: the cyber operations context. In Czosseck, R. Ottis, & K. Ziolkowski (Eds.), *Fourth International Conference on Cyber Conflict*. Tallinn: NATO CCD COE Publications.
Schmitt, M. (2015). *Foreword to Jens David Ohlin, Kevin Govern, and Claire Finkelstein (eds) Cyberwar: law and ethics for virtual conflicts*. Oxford: Oxford University Press.
Schmitt, M., & Vihul, L. (2016). The emergence of international legal norms for cyberconflict. In F. Allhoff, A. Henschke, & B. J. Strawser (Eds.), *Binary bullets: the ethics of cyberwarfare*. New York: Oxford University Press.
Stone, J. (2013). Cyber war will take place! *Journal of Strategic Studies, 36*(1), 101–108.
Taddeo, M. (2014). What ethics has to do with the regulation of cyberwarfare. *Ethics and Armed Forces, 2*, 36–40.
Van der Linden, H (2012). 'On the violence of systemic violence: A critique of Slavoj Žižek', *Radical Philosophy Review, 15*:33–51.
Walzer, M. (1977). *Just and unjust wars: a moral argument with historical illustrations*. New York: Basic Books.
Wolff, R. P. (1969). On violence. *Journal of Philosophy, 66*(19), 616–628.