Review of International Studies, Vol. 44, part 2, pp. 324–342. doi:10.1017/S026021051700047X
© British International Studies Association 2017
First published online 17 October 2017

RIS

Just cyber war?: *Casus belli*, information ethics, and the human perspective

Matt Sleat*
University of Sheffield

Abstract

Does the advent of cyber war require us to abandon the traditional ethical framework for thinking about the morality of warfare – just war theory – and develop principles specific to the unique nature of cyber attacks? Or can just war theory still provide an appropriate basis for thinking through the ethical issues raised by cyber weapons? This article explores these questions via the issue of whether a cyber attack can constitute a *casus belli*. The first half of the article critically engages with recent attempts to provide a new theory of just information warfare (JIW) that is supposedly better suited to the unique character of cyber war insofar as it is grounded the broader meta-ethical framework of information ethics (IE). Yet the article argues that not only is JIW fundamentally unsuitable as a way of thinking about cyber war, but (in the second half) that it is possible to develop a different account of how we can understand a cyber attack as constituting *a casus belli* in a way that is in keeping with traditional just war theory. In short, there is no need to reinvent just war theory for the digital age.

Keywords

Cyber War; Just War Theory; *Casus Belli*; Information Ethics; Violence

The bewildering pace of technological change has transformed what looked like science fiction alarmism twenty years ago – the offensive use of computer technology to target other computer systems, infrastructures, networks, or personal devices – and raised it rapidly into a position of significant prominence in security and policymaking circles. The discovery of the so-called Stuxnet worm in 2010, which destroyed several centrifuges at Iran's uranium enrichment facility in Natanz, the use of cyber attacks during the tensions between Russia and Georgia and Estonia in 2008 and 2007 respectively, and several other recent cyber incidents, have meant that anyone reflecting on the future of warfare in the twenty-first century needs to take seriously the presence and use of cyber weapons.[1] The novelty of these weapons poses the question of whether the advent of cyber war require us to abandon the traditional ethical framework for thinking about the morality of warfare – just war theory – and develop principles specific to the unique nature of cyber attacks. Or can existing just war theory, maybe suitably amended but without fundamental alteration, still provide an appropriate basis for thinking through the ethical issues raised by cyber weapons?

We can identify three broad schools of thought in response to this question. The first is sceptical either that the development of cyber weapons raises any novel ethical questions that are not already covered by existing just war theory or that cyber attacks ought to be understood under the paradigm

* Correspondence to: Matt Sleat, Department of Politics, Elmfield, Northumberland Road, Sheffield, S10 2TU. Author's email: m.sleat@sheffield.ac.uk

[1] For a good survey of cyber conflict including Stuxnet and Russia's conflict with Georgia and Estonia, see Jason Healey (ed.), *A Fierce Domain: Conflict in Cyberspace, 1986–2012* (Cyber Conflict Studies Association, 2013).

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

of warfare in the first place. These 'sceptics' need not deny that cyber war represents a significant shift in contemporary warfare, they only dispute whether these developments are such that they raise any new ethical questions.[2] The second school of thought consists of those we may call 'moderates' who accept that cyber war creates some interesting novel moral, philosophical, and legal questions, but believe that just war theory contains resources that can be expanded or amended in order to capture the ethical issues that the innovations of cyber war have generated.[3] The final school, that of the 'radicals', offers a somewhat more far-reaching philosophical and ethical response. It believes that the questions posed by cyber war cannot be adequately addressed by traditional just war theory because this new 'domain' of warfare is so radically different from all forms of conflict which came before it that we find ourselves with a 'regulatory gap' that can only be bridged via a new meta-ethical framework of analysis.[4] The old one, suitable though it may have been for nuclear weapons, tanks, jet fighters, and the like, is simply not appropriate for the age of cyber war. In the words of John Arquilla, in many ways the first person to anticipate the emergence of cyber warfare and its ethical implications, cyber war leaves 'just war theory in tatters'.[5]

There is something admittedly paradoxical-sounding about the moderate position and intuitive in that of the radicals. After all, the advent of cyber weapons surely represents a revolutionary development in our security environment and hence we should not expect ethical frameworks such as just war theory, developed in and for a pre-digital world, to apply in this new context. Yet that is what I wish to argue here. And I wish to do so via the *jus ad bellum* principle of just cause, which has been one of the central issues around which these debates have hitherto focused. If as Vitoria put it, 'There is a single and only just cause for commencing a war, namely, a wrong received', does a cyber

[2] Roger Crisp, 'Cyberwarfare: No New Ethics Needed', available at: {http://blog.practicalethics.ox.ac.uk/2012/06/cyberwarfare-no-new-ethics-needed/} accessed 14 August 2017; Eric Gartzke, 'The myth of cyberwar: Bringing war in cyberspace back down to Earth', *International Security*, 38:2 (2013), pp. 41–73; Larry May, 'The nature of war and the idea of "cyberwar"', in Jens David Ohlin et al. (eds), *Cyberwar – Law and Ethics for Virtual Conflicts* (Oxford: Oxford University Press, 2015), pp. 3–15; Thomas Rid, *Cyber War Will Not Take Place* (Hurst & Co.: London, 2013); Thomas Rid, 'Cyber war will not take place', *Journal of Strategic Studies*, 35:1 (2012), pp. 5–32.

[3] See particularly George Lucas, *Ethics and Cyber Warfare: The Quest for Responsible Security in the Age of Digital Warfare* (Oxford: Oxford University Press, 2017). See also Edward T. Barrett, 'Reliable old wineskins: the applicability of the just war tradition to military cyber operations', *Philosophy & Technology*, 28:3 (2015), pp. 387–405; Edward T. Barrett, 'Warfare in a new domain: the ethics of military cyber-operations', *Journal of Military Ethics*, 12:1 (2013), pp. 4–17; James Cook, 'Is there anything morally special about cyberwar?', in Ohlin et al. (eds), *Cyberwar – Law and Ethics for Virtual Conflicts*, pp. 16–36; James Cook, '"Cyberation" and just war doctrine: a response to Randall Dipert', *Journal of Military Ethics*, 9:4 (2010), pp. 411–23; Dorothy E. Denning and Bradley J. Strawser, 'Moral cyber weapons', in Luciano Floridi and Mariarosaria Taddeo (eds), *The Ethics of Information Warfare* (Springer: London, 2014), pp. 85–103; Christopher J. Eberle, 'Just war and cyberwar', *Journal of Military Ethics*, 12:1 (2013), pp. 54–67; Ryan Jenkins, 'Is Stuxnet physical? Does it matter?', *Journal of Military Ethics*, 12:1 (2013), pp. 68–79. The moderate school has its legal correlative, such as the authors of the *Tallinn Manual*, which takes existing international law and norms surrounding *jus ad bellum* and *jus in bello* apply to cyber operations. Michael N. Schmitt et al., *Tallinn Manual on the International Law Applicable to Cyber Warfare* (Cambridge: Cambridge University Press, 2013).

[4] Ludovica Glorioso, 'Cyber conflicts: Addressing the regulatory gap', *Philosophy & Technology*, 28:3 (2015), pp. 333–8.

[5] John Arquilla, 'Ethics and information in warfare', in Z. Khalizad et al. (eds), *The Changing Role of Information in Warfare* (Santa Monica: Rand Corporation, 1999), p. 394. See also John Arquilla, 'Twenty years of cyberwar', *Journal of Military Ethics*, 12:1 (2013), pp. 80–7; Selmer Bringsjord and John Licato, 'By disanalogy, cyberwarfare is utterly new', *Philosophy & Technology*, 28:3 (2015), pp. 339–58; Randall Dipert, 'The ethics of cyberwarfare', *Journal of Military Ethics*, 9:4 (2010), pp. 384–410.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

attack constitute a 'wrong received'?[6] In particular, can a cyber attack properly be understood as an aggressive act? And how does our answer to that question then impinge on whether the use of a cyber weapon constitutes a *casus belli*? The sceptics tend to reject the notion that cyber attacks can be acts of aggression, preferring to see them as analogous to other acts that fall short of aggression such as embargoes or espionage.[7] And if cyber attacks are not a form of aggression then they cannot be an act of war, and certainly not ones that are to be regulated either by the principles of just war theory or of international law (such as UN Charter Articles 2.4 or 51). No aggression, no *casus belli*. Most commentators, however, fall into one of the other two camps, believing either that cyber attacks can be understood as aggressive in terms that are straightforwardly consistent with the terms of just war theory (the moderates) or that cyber attacks are indeed aggressive but that to understand how and why we have to adopt a radically different meta-ethical framework (the radicals).

The strategy of this article is to demonstrate the sufficiency of traditional just war theory in two stages. In the first half of the article I analyse one of the most sophisticated and influential branches of the radical school approach that suggests just war theory needs to be 'merged' with a new and more pertinent meta-ethics. The information age needs an information ethics (IE), and it is only by rethinking our most basic ontological and moral assumptions along these lines that we can develop an account of just war theory – just information war (JIW) – appropriate for regulating the ethical use of cyber weapons.[8] I shall argue that JIW not only delivers inadequate answers to the questions that we need a theory of just war applied to cyber attacks to provide, more troublingly it is also unable to get a sufficient grasp on the questions that a theory should enable us *to ask*. Despite the professed novelties of cyber attacks and the temptation to think that conflict in the cyber realm might demand a new meta-ethics designed specifically to incorporate developments in ICTs, the lesson from this discussion will be that any adequate ethical theory for regulating even conflict in the cyber realm needs to retain what I shall call the 'human perspective'. In the second half of the article I set out an account of what such a perspective could look like. I seek to demonstrate that of those ontological features of cyber attacks radicals take to be incompatible with traditional just war theory – that they are *non-physical*, target *non-humans* and are *non-violent* – only the latter actually poses any sort of meaningful challenge. The remaining two are either descriptively dubious or turn out to be of little relevance to just war theory. Yet even the challenge of non-violence can be met quite straightforwardly by conceptualising the harm caused by cyber attacks in terms of harm caused to vital human interests through degrading the functionality of computer systems necessary to a country's critical infrastructure. And as such it is possible to conceive the possibility that even non-violent cyber attacks (for not all cyber attacks are non-violent) may, under certain conditions, represent a *casus belli* in a way that is largely in keeping with traditional just war theory. In short, there is no need to reinvent the just war wheel for the digital age.

Information ethics and just war theory

It is a central contention of IE that addressing cyber attacks 'solely on the basis of JWT [just war theory] generates more ethical conundrums than it solves'.[9] This is because JWT makes a series of ontological assumptions about the nature of war that are fundamentally incompatible with the

[6] Cited in Michael Walzer, *Just and Unjust Wars* (4th edn, Basic Group: New York, 2006), p. 62.
[7] May, 'The nature of war and the idea of "cyberwar"'; Rid, *Cyber War Will Not Take Place*; Rid, 'Cyber war will not take place'.
[8] Mariarosaria Taddeo, 'Just information warfare', *Topoi*, 35:1 (2016), pp. 213–24; Mariarosaria Taddeo, 'Information warfare and just war theory', in Floridi and Taddeo (eds), *The Ethics of Information Warfare*, pp. 123–38.
[9] Taddeo, 'Just information warfare', p. 216.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

character of cyber attacks. For our purposes we can identify three related characteristics of cyber attacks that can then each be used to identify corresponding shortcomings with JWT:[10]

Non-physical – JWT focuses mainly on the use of (usually kinetic) force in the domain of physical objects. As Brian Orend puts it, 'the gold standard of casus belli is a kinetic physical attack'.[11] Cyber attacks take place in a non-physical or virtual domain involving non-physical or virtual entities (though they may have consequences for physical objects in the physical realm, which is the sense in which they can be transversal).[12]

Non-human – JWT rests on an 'anthropocentric ontology' that prioritises respect for humans and their interests, rights, etc., and disregards all non-human entities. The target of cyber attacks are computer systems and the information or data contained within them.[13]

Non-violent – JWT understands violence primarily in terms of injury caused to human beings or the destruction of physical objects. Yet not only is the harm that cyber attacks cause often to non-physical, non-human entities (see above) but it is also best conceived in terms of the (mal)*functioning* of computer systems, which often leads to neither bloodshed nor physical destruction. They may include, for instance, the use of a virus or a DDoS attack capable of disrupting or denying enemy access to information which may cause severe damage to the opponent yet do not appear to resemble traditional forms of violence.[14] Such losses of computer functionality will often be temporary and reversible following a cyber attack (for example, full system functionality will likely be restored as soon as a DDoS attack is halted).

We can understand these characteristics as identifying the domain in which cyber attacks take place, the status of the entities involved in them, and the nature of the harm caused to those entities. Where JWT is concerned with violent acts (harm) against human beings (target) in the physical realm (domain), cyber attacks are non-violent acts against non-human entities in a non-physical realm. Hence any appropriate ethical framework for thinking about cyber attacks must be able to incorporate their ontological particularities in a way that JWT supposedly cannot.

The motivation for developing an ethics of information stems from the broader thought that ICTs in general 'are re-ontologising the context in which ethical issues arise, and in doing so they not only transform old problems, but also invite us to explore afresh the foundations on which our ethical positions are based'.[15] ICTs have fundamentally transformed the intrinsic nature of reality, blurring the divide between offline and online, notions of personal identity, space, time, and history, as well as conceptions of moral agency and responsibility, and, crucially, the very criterion for existence (being). It is no longer something's immutability, nor its being potentially subject to perception that

[10] Ibid.; Dipert, 'The ethics of cyberwarfare'.
[11] Brian Orend, *The Morality of War* (2nd edn, Ontario: Broadview Press, 2013), p. 176.
[12] To be clear from the outset, one of the arguments I am going to make against JIW is that ontological questions regarding whether or not cyber attacks are physical are simply irrelevant to assessing whether they can constitute a *casus belli*. Hence, while there is a live debate about the issue of cyber attacks' (non-)physicality, which are of interest in their own right, they nevertheless do not undermine traditional JWT in the way advocates of the radical approach suggest. See Jenkins, 'Is Stuxnet physical?'; George R. Lucas, 'Postmodern war', *Journal of Military Ethics*, 9:4 (2010), pp. 289–98.
[13] See also May, 'The nature of war and the idea of "cyberwar"'.
[14] See also Arquilla, 'Ethics and information in warfare'; Barrett, 'Warfare in a new domain'; Dipert, 'The ethics of cyberwarfare'.
[15] Luciano Floridi, *The Ethics of Information* (Oxford: Oxford University Press, 2013), p. 56.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

qualifies it for existence, but its being something with which we can interact. 'To be is to be interactable.'[16] And it is no longer only material objects with which we, as physical beings, can interact, but also with non-physical entities such as programs, databases, music files, virtual avatars and assets, and so on, all of which are forms of information.

The trouble with our standard ethical theories is that they employ a level of analysis (LoA) that excludes these non-physical but nevertheless real entities.[17] A LoA is a theoretical tool that we employ in order to focus on some aspects of a system and to ignore others, depending on the interests or goals of the observer. Imagine we are analysing a car: an engineer will likely focus on the car's aerodynamics, the efficiency of its parts, weight, material, and so on. A potential buyer of the same car may focus on its aesthetic appearance. Both the engineer and the buyer are focusing on the same car but they endorse different LoAs in order to pick out particular features of the system.[18] The higher the LoA, the more features of the system it will incorporate. The lower the LoA, the less observables it will include. But which LoA we adopt, which perspective we employ, will depend on what features of the system are relevant to the purposes of any particular observation.

Many of our standard ethical theories employ androcentric, anthropocentric, or biocentric LoAs, for example, and in doing so are then unable to support an effective analysis of moral scenarios in which non-human informational entities are involved. Even atypical ethics, like animal or environmental ethics, are still biocentric and hence biased against what is inanimate, lifeless, intangible, abstract, engineered, artificial, synthetic, or merely possible.[19] They ascribe moral concern or value only to what we intuitively think of as alive or able to feel pain, and hence cannot account for the moral scenarios that involve non-human entities like corpses, historical artefacts, or works of art, or non-physical entities like past and future generations, ideas, software, viruses, and data. From an anthropocentric LoA such entities can only have instrumental value at best. IE is committed to a LoA that interprets reality informationally. And because everything that exists (including physical objects) can be understood in terms of being a discrete, self-contained, encapsulated package containing appropriate data structures (which constitute the nature of the entity) and a collection of operations, functions, or procedures activated by various interactions or stimuli, it is a LoA that encompasses and concerns the entire realm of reality – now redescribed as the *infosphere* in order to emphasise its essentially informational nature.[20] This is the highest possible LoA because it seeks to include all entities that meet the minimum common characteristic to exist – being informational – and in doing so becomes synonymous with reality or Being itself.[21]

It is in this sense that IE is *ontocentric*: it focuses on what exists, not on the human or the biological (though it will include these insofar as they are informational entities). And it applies a principle of ontological equality to all existing entities (physical and non-physical) in which

> every informational entity, insofar as it is an expression of Being, has a dignity constituted by its mode of existence and essence … This dignity *prima facie* deserves to be respected and

[16] Ibid., p. 10.
[17] Ibid., ch. 3; Luciano Floridi, 'Information ethics: its nature and scope', in Jeroen Van Den Hoven and John Weckert (eds), *Information Technology and Moral Philosophy* (Cambridge: Cambridge University Press, 2010), pp. 40–65; Taddeo, 'Just information warfare'.
[18] Taddeo, 'Just information warfare', p. 215.
[19] Floridi, *The Ethics of Information*, p. 64.
[20] Floridi, 'Information ethics', p. 46.
[21] Floridi, *The Ethics of Information*, p. 6.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

> hence may place moral claims on any interacting agent. It ought to contribute towards constraining and guiding her ethical decisions and behaviour, even if only initially and in an overridable way. This ontological equality principle means that any form of reality – that is … any instance of information – simply for the fact of being what it is, enjoys an initial, overridable, minimal right to exist and develop in a way appropriate to its nature.[22]

Once we adopt the information approach the entire universe and everything in it – physical and non-physical – becomes a potential centre of moral concern simply by virtue of their informational nature ('something more elementary and fundamental than life and pain').[23] If something exists, then it has a dignity that deserves respect.

This respect is substantiated in practice by recognising that any informational entity has a right to persist in its own status, that is, to continue being *that* sort of information,[24] and a right to flourish through improving and enriching its existence and essence.[25] IE is *patient-orientated* insofar as it considers the morality of any action in relation to its effects on the particular informational entity that is the recipient of that action.[26] More generally, however, moral approval or disapproval should also be based on how an action affects the well-being (the enrichment or impoverishment) of the whole infosphere.[27] This is conceptualised in terms of entropy, though this is not to be thought of in the thermodynamic sense of the state of randomness or disorder in a physical system at the atomic level, but as indicating 'the decrease or decay of information leading to absence of form, pattern, differentiation, or content in the infosphere', specifically in terms either of destruction ('the complete annihilation of the entity in question, which ceases to exist') or corruption ('a form of pollution or depletion of some of the properties of the entity, which ceases to exist as that entity and begins to exist as a different entity minus that properties that have been corrupted or eliminated').[28] This notion of entropy enables IE to determine what is morally right or wrong with reference to four principles:[29]

0. entropy ought not to be caused in the infosphere (null law)
1. entropy ought to be prevented in the infosphere
2. entropy ought to be removed from the infosphere
3. the flourishing of informational entities as well as the whole infosphere ought to be promoted by preserving, cultivating, and enriching their properties

That an action is to be judged according to how it affects the flourishing of the infosphere, with the blooming of the infosphere representing the ultimate good and its corruption or destruction the ultimate evil, makes IE an *ecological* form of meta-ethics.

As a meta-ethical position, IE can be and has been applied to an array of ethical issues, not all of which need to be related to ICTs, in order to generate moral understanding and practical guidance.

[22] Ibid., p. 69
[23] Ibid.
[24] Massimo Durante, 'Violence, just cyber war and information', *Philosophy & Technology*, 28:3 (2015), pp. 369–85.
[25] Floridi, 'Information ethics', p. 48.
[26] Taddeo, 'Just information warfare', p. 220.
[27] Floridi, *The Ethics of Information*, p. 70.
[28] Ibid., p 67.
[29] Ibid., p. 71; Floridi, 'Information ethics', pp. 58–9; Taddeo, 'Just information warfare', p. 221; Taddeo, 'Information warfare and just war theory', p. 133.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

When it comes to the question of the wrong caused by cyber attacks, IE answers that all information systems are possible moral patients that can suffer harm, and the morality of a cyber attack is to be assessed on the basis of its effects on the right of those entities to exist and flourish, as well as the general flourishing of the infosphere. Cyber attacks therefore fail to respect information's deserved dignity in being. Any entity (for example, hacker or computer virus) that comes into conflict (causes entropy) with other entities such as software or an information system may lose its right to exist. Indeed, because evil is defined as an increase in the level of entropy within the infosphere, it 'is the moral duty of the other inhabitants to remove such a malicious entity from the environment or at least impede it from perpetrating more evil'.[30] From this we can derive the following condition of an information *casus belli*:

I. Cyber war ought to be waged only against those entities that endanger or disrupt the well-being of the infosphere

And then two further principles of *jus ad bellum*:

II. Cyber war ought to be waged to preserve the well-being of the infosphere

III. Cyber war ought not to be waged to promote the well-being of the infosphere.[31]

The second principle limits the objectives of cyber war to repairing the damage done to the infosphere to the level prior to when the malicious entity began to increase entropy within it; the promotion of the well-being of the infosphere lies beyond the remit of a just cyber war (the third principle). In this sense we might think of these two principles as an extension of the condition of right intention.

Against just information war theory

The advantage of JIW, according to its advocates, is that it is a meta-ethic with ontological assumptions consistent with the peculiar character of cyber attacks, in contrast to traditional JWT. To recall, JWT is concerned with violent acts (harm) against human beings (target) in the physical realm (domain), whereas cyber attacks are non-violent acts against non-human entities in a non-physical realm. As we have seen, IE is a *patient-oriented*, *ontocentric*, and *ecological* meta-ethics.[32] That IE is ontocentric means that it accounts for entities and actions that exist in non-physical and physical domains (and those, like cyber attacks, which might transverse the two). That IE is patient-oriented allows us to make sense of how non-human entities, including information, can be the subject of ethical concern. And that IE is an ecological meta-ethics provides a way of conceptualising the harm caused to all information entities. As such, IE provides a novel answer to the question posed earlier in the article as to the nature of the wrong involved in a cyber attack: such actions harm the information entities that it targets, and in doing cause entropy in the infosphere.[33] This explains both the nature of their 'wrongness' and, when worked up into an account of JIW, the grounds on which a cyber attack can constitute a *casus belli*. It is only legitimate to engage in cyber war when the infosphere is endangered or disrupted. This reaffirms the notion that a just war will be reactive, responding to prior acts and cannot itself be the first act of aggression, which is in keeping with the traditional just war theory notion of just cause (that the notion of 'endangering' is included in the principle would seem to legitimate pre-emptive attacks also, though I leave that to one side here).

[30] Taddeo, 'Just information warfare', p. 221.
[31] Ibid., p. 221; Taddeo, 'Information warfare and just war theory', p. 134.
[32] Floridi, *The Ethics of Information*; Taddeo, 'Just information warfare'.
[33] Taddeo, 'Information warfare and just war theory', p. 136.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

But this cannot be the whole story; we need more from an ethical framework for thinking about cyber warfare than this. Why? Because if we leave matters here then we are left with a theory that tells us only, in the words of one JIW advocate, that a disruptive act that damages, deteriorates, deletes, or suppresses an information entity 'patently rises to the level of warfare (for instance, when destroying or damaging a computing infrastructure)'.[34] *Any* disruption to *any* informational entity is an act of aggression? This is either a dangerous or an absurd notion, likely both. But it is, for reasons I now want to set out, the position that a theory of just war merged with IE will necessarily be committed to.

Remember that *all* informational entities, by virtue of being, have a *prima facie* right to exist. All things being equal, a world in which a particular email or rock exists is preferable, even if only slightly so, to a world that is the same in all ways apart from the absence of that same email or rock because that world will be ontologically richer.[35] The same will be true of all and any other informational entities. Likewise, all things again being equal, from the perspective of IE there is no relevant moral difference between different informational entities at the level of their right to exist (unless they are disrupting the infosphere, in which case they forfeit that right). This is the principle of ontological equality central to IE. So the loss of or disruption to any informational entities would seem to be an equal wrong. But this generates some very disturbing conclusions. Imagine a cyber attack that took the form of a virus that trawled through every computer network in the world, public and private, and systematically deleted every 'selfie' it can find. These pictures are information entities that have a dignity due to their existence that the creator of the virus has failed to respect. The number of selfies in existence is probably easily in the billions and hence the infosphere will be a significantly diminished place because of this virus. Now imagine a different virus that sought not to delete selfies but to delete or corrupt a single program, though one that is absolutely essential to the functioning of a country's major power grid. JIW tells us that the use of either virus would constitute a *casus belli*.

If that conclusion looks problematic, as it should, it is not difficult to see why. There are other considerations that are essential to the question of whether something represents an act of aggression or not, none of which JIW provides us with the resources to judge in any particular scenario. The first is the target of the attack. Though the alternative scenarios set up above purposefully differ in their severity, they draw attention to the fact that when assessing whether an attack represents a *casus belli* we need to be able to discriminate between the nature, value, and significance of the target of those attacks in a way that is not possible from an informational perspective. *Any* action that diminishes the well-being of the infosphere makes the offender a legitimate target of a violent response. JIW does not *and cannot because of the principle of ontological equality* distinguish between different informational entities, a consequence of which must be that any cyber attack on any entity will unconditionally represent a just cause. But that is obviously outrageously permissive. We think, and rightly so, that there are some entities even the destruction of which would not justify a violent response; indeed we probably think the vast majority of entities that exist (physical and non-physical) would fall into this category. But such judgements cannot be made from the informational LoA where the principle of ontological equality gives all entities the same *prima facie* right to existence. So the mere 'disruption of the infosphere' cannot in itself constitute an act of aggression. And to think that it does is to offer no guidance as to when it is legitimate to go to war other than to say that only an aggressive act can provide the grounds for a just conflict. That much we already know.[36]

[34] Durante, 'Violence, just cyber war and information', p. 371.
[35] Floridi, *The Ethics of Information*, p. 131.
[36] Another way of thinking about this deficiency with the IE account is via the condition of proportionality rather than just cause, as we are doing here. We might think that the reason why the virus that deletes selfies does not

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

Even if we can identify a category of informational entities the targeting of which represents an act of aggression, we are still required to make judgements as to how disruptive or grave an attack is before we can judge whether it represents a *casus belli*. JIW tells us only that the disruption of the infosphere is sufficient cause for a just war. It says nothing about how significant that disruption has to be, regardless of what the target is. Surely there are some attacks, even on vital targets such as critical infrastructure, which are not significant enough to warrant or legitimate a violent response. DDoS attacks such as those launched against Estonia in 2007 did target critical infrastructure, but the consensus (rightly) seems to be that though they were disruptive they nevertheless fell short of being aggressive, and are better characterised as either vandalism or criminal activities rather than acts of war.[37] To be sure, and as Stuxnet has proven, cyber attacks *can* destroy objects in the physical world. We may find that these sort of destructive attacks are the exception more than the rule, yet they are possible, and the authors of the *Tallinn Manual* (largely) agreed that the Stuxnet worm did represent a *casus belli*; Iran would have been justified in retaliating.[38] Furthermore, it is not unreasonable to assume that a cyber attack may directly lead to human casualties in the future. Failure to take these considerations into account leaves us unable to distinguish between different categories of action in the cyber realm which, as Lucas rightly points out, is an unfortunate trend in contemporary discussions. Too little attempt is often made to distinguish between cyber vandalism, cyber crime, cyber espionage, cyber terrorism, and cyber war,[39] and part of how we make such distinctions is going to rely upon judgements as to the perceived severity of any harm caused. Questions of whether a cyber attack is an annoyance or a menace, an aggressive act or act of vandalism, are crucial for determining whether it can constitute *a casus belli* or not. And it is far from clear how JIW can allow us to make these judgements, again because of the sort of meta-ethical approach that IE is. Whereas the principle of ontological equality undermines our ability to discriminate between targets, as informational entities have an equal *prima facie* right to existence, when that is combined with the commitment to being a patient-oriented meta-ethics then that necessarily draws our attention away from questions that we might otherwise deem morally pertinent to focus *exclusively* on the wrong done to *that* entity/patient. To adopt any other perspective, such as the indirect consequences of an attack on humans and their interests, is to have failed to treat the target entity as the victim of a moral wrong.

This relates to a third consideration, which in a certain sense ties the other two together also, and that is the apolitical nature of the JIW account. From the perspective of the infosphere what we see are only ontologically discrete informational entities. Hence it follows quite straightforwardly from JIW being an environmental meta-ethic that *any* entity that disrupts or endangers the well-being of the infosphere becomes a licit target, and that it becomes a moral duty for *all* other entities to prevent that licit entity from causing more evil.[40] What JIW portrays is a world in which any entity can be declared war upon, properly speaking, and any entity can declare war. But this totally misses the necessarily political dimension of war. War is a relationship between groups of human beings, and specifically between political groups. We primarily think of these groups in terms of states, but the spread of irregular forms of warfare over recent decades that involve non-state actor complicates

justify a violent response is because the principle of proportionality tells us that no armed attack could ever be a fitting response to such an inane act. That seems obviously right, but again it is unclear how this assessment can be made from the informational perspective as it requires us to make judgements as to the value and significance of the entities damaged in the attack, which we can only do from a perspective other than of the infosphere.

[37] Lucas, *Ethics and Cyber Warfare*, p. 117; Schmitt et al., *Tallinn Manual*, p. 75.

[38] Schmitt et al., *Tallinn Manual*, p. 58. See also Lucas, *Ethics and Cyber Warfare*, pp. 117–18.

[39] George R. Lucas, 'Permissible preventative cyberwar: Restricting cyber conflict to justified military targets', in Floridi and Taddeo (eds), *The Ethics of Information Warfare*, pp. 73–83.

[40] Taddeo, 'Information warfare and just war theory', p. 136.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

that picture. Nevertheless even insurgents or terrorists are understood as political in the sense of having and pursuing specifically political objectives through their actions. War is, in Clausewitz's famous remark, 'a mere continuation of policy by other means'. This matters because only an aggressive action by a political group can constitute a *casus belli*. A group pursuing financial gain via illegal cyber activities, for instance, could be the subject of legal proceedings, but that is not war (and it is important that it is not war). The same would be true of an individual hacker, even if he might be acting on or pursuing political ends (without direction from any government). Any perspective that does not appreciate the specifically political dimension of warfare is therefore an inappropriate place from which to view the ethical problems of cyber war.[41]

All of this is supposed to lend credence to the thought that causing harm to the infosphere cannot, in itself, constitute a *casus belli*. It is not the case that any cyber attack will represent a just cause. In order to make appropriate judgements as to whether a cyber attack does constitute a *casus belli* there is no escaping the fact that we can only do so from a specifically human perspective or, to use the discourse of IE, a human LoA. This is because a judgement on the justness of a war is a judgement of human actions in relation to human interests and projects and with regard to future human consequences. While it might be possible to recast cyber attacks in terms of their entropic effects on the infosphere, it seems fairly clear that the judgements we can and do make, *relatively unproblematically*, about the wrong of cyber attacks relates directly to their effects on human interests. To provide an answer to the question of whether an act is aggressive or not we need to assess the content of the wrong inflicted by one political party on another, which itself requires accompanying judgements about the nature and relative importance of a large and complex set of specifically *human* needs, interests, values and purposes, as well as the extent of the harm or disruption that that act does to them. The justice of a war is a question asked by humans of specific human actions and their effects on a particular group of other human beings. A theory of just war needs to speak directly to that human question.

Stating that defending the well-being of the infosphere is a just reason to go to war is not an answer to the question we need an answer to. We need to know precisely *which* bits of the infosphere can justly attack *which* other bits, which political entities can justly wage war against which other political entities. And that means we need to identify quite clearly exactly what the nature of the wrong was, by whom to whom. To do that, the perspective from which we need to judge, assess, critique, and ultimately act is not and cannot be the perspective of the infosphere and its well-being, but our local and necessarily more limited human perspective of political groups, their actions, and their impact on our projects and purposes. War is by its very nature both a human activity and one that divides humans against each other, the purpose of which is to work out competing human claims through violence (to resources, values, interests, etc.). Judgements as to the justness of that violence must therefore remain tied to the human condition from which it arises. They cannot be, and we do not need them to be, judgements from the perspective of their effects on the infosphere.

IE: Minimalism not reductionism? Or, in favour of human ethics

At this stage I may legitimately be accused of having mistaken IE's endorsement of a minimalist approach, insofar as it considers informational nature as the minimal common denominator among all existing things, for a crude reductionism: IE 'does not claim that the informational approach is the

[41] That JIW states we have a *moral obligation* to attack licit entities, rather than simply the right, is also a huge difference between it and traditional just war theory that would have momentous legal ramifications, and deserves much more attention than it is given in the literature.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

unique LoA from which moral discourse is addressed. Rather, it maintains that the informational LoA provides a minimal starting point, which can then be enriched by considering other moral perspectives.'[42] What IE tells us is that in considering ethical issues we must adopt whichever LoA is most appropriate, and that need not always be the highest level of abstraction, that is, the informational perspective. For some issues, an anthropocentric LoA may indeed be more suitable. So, for instance, it is worth remembering that IE insists that all informational entities only have the *prima facie* right to exist, a right that can be overridden one assumes if we decide from whichever LoA we do adopt that particular entities are of greater moral significance. But if this is right it either makes IE redundant or uninteresting as a way of approaching cyber war. Or, put differently, my claim is not that IE is mistaken in taking the informational approach to be the 'unique' LoA from which to assess moral scenarios, but that it does not represent even a relevant 'minimal starting point' for making judgements of the sort that the ethical questions relating to the wrong of cyber attacks require.

There are two ways of getting at this point. The first is to insist that there is nothing illegitimate or ethically inappropriate about starting our moral reasoning directly from an anthropocentric LoA, that doing so does not represent some arbitrary bias or unjustified prejudice. This is to register a fairly fundamental disagreement with IE as a meta-ethic. The second is to deny that the informational LoA adds any morally relevant information to our reasoning regarding cyber war; everything that is needed to explain the wrong of cyber attacks is readily at hand in an anthropocentric LoA. The anthropocentric LoA is *sufficient* for our moral purposes. While these are strictly speaking independent claims insofar as neither implies or relies upon the other, I want to try to defend them both via a similar route.

In defending the move to an ontocentric LoA, Floridi tells us: 'It seems that any attempt to exclude nonliving entities is based on some specific, low LoA and its corresponding observables, but that *this is an arbitrary choice*. In the scale of beings, there may be no good reasons to stop anywhere but at the bottom … There seems to be no good reason not to adopt a higher and more inclusive, ontocentric LoA.'[43] Everything here turns on what counts as a 'good reason'. Why not settle at the lower anthropocentric LoA? It is certainly the case that the thought that human beings occupy a position of *absolute* importance in the universe is no longer a plausible one that most of us can truthfully hold. But it is an error to think that because we cannot stop at anthropocentrism due to our moral significance from the point of view of the cosmos that we must therefore go straight to the 'bottom' and start our moral reasoning from the minimal condition of being, that to do otherwise would be simply arbitrary. The mistake is to think that there are only two ethical positions available to us: either human beings are of moral significance from some non-human perspective or they have no greater moral significance than any other entity that exists. While human beings, their interests, values, and projects are of no particular interest to the universe (or the infosphere), they are certainly of supreme significance *to us*. They are important from our perspective. Of course there is little surprising about that, after all it is human beings who ask ethical questions of their own activities, and it is with other human beings that we seek to discuss and explore these issues. Even when the question relates to our behaviour towards other non-human entities, it is still us who ask the questions and us who judge the adequacy of the answers. But that our own moral significance is a central feature of our ethical experience, and we know of no other (and even if we did, it is far from obvious how we would or should relate to it), is certainly *a* reason to think that it should be taken

[42] Taddeo, 'Information warfare and just war theory', p. 132.
[43] Floridi, 'Information ethics', p. 56, emphasis added.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

seriously in our moral thinking. To deny that this is a 'good enough' reason is to hanker for the sort of ultimate normative justification from some non-human perspective (God, Nature, Reason, Infosphere) that is no longer available to us, if it ever was.[44]

It is humans, and only humans, to whom ethical justifications for human acts have to be offered. The moral discourse we need is therefore one in which the features that *we* deem to be ethically relevant to assessing a cyber attack are readily at hand. This is precisely why traditional just war theory has tended towards a very minimal notion of what constitutes a *casus belli*. Whatever else we may disagree on, we readily recognise the harm caused by violent acts, and even if we think they are justified it is important that we acknowledge how they stand in need of justification in the first place. The difficulty with cyber attacks is that it is unclear if we can understand them as aggressive at all, as was noted earlier. That they are non-violent acts against non-human entities in a non-physical realm were perceived as reasons for needing to shift to a new meta-ethical framework in which cyber attacks can be understood as acts of aggression. But if we can show how there are sufficient resources for understanding cyber attacks as aggressive acts from within an anthropocentric LoA then any further ethical light that is shone on the issue from an informational perspective is ultimately superfluous, and hence the informational LoA does not provide even a 'minimal starting point' for our ethical reflection on cyber war. In fact, when it comes to the question of whether a cyber attack represents a just cause it can add nothing relevant to our reasoning that is not already present in the human perspective. Or so I now need to demonstrate.

Cyber war from a human perspective

As we have seen, advocates of IE insist that cyber attacks can only be included in a just war perspective if we adopt a new meta-ethical framework that endorses ontological commitments that can allow us to make sense of non-physical acts of non-violence against non-humans as nevertheless aggressive. So far I have argued that IE fails to provide an adequate framework for thinking about the possibility of a cyber *casus belli*. Yet that does not in itself prove that the 'radicals' are wrong in thinking that there is a fundamental incompatibility between cyber attacks and JWT in the first place. They may still be right about that but only wrong in how they develop their alternative account of JIW. Hence if I want to defend a traditional JWT approach to thinking about cyber attacks then I need to show how those peculiarities of conflict in the cyber realm – that they are *non-violent* and *non-physical* attacks targeting *non-humans* – can be squared with its basic ontological assumptions. This is what I wish to do in this final section.

The issue of cyber attacks' non-physicality can be dealt with fairly quickly. While there is something intuitive in the thought that the cyber domain is not equivalent to the other domains of warfare – sea, air, land, and space – this is not adequately captured by the sense that this is because it is a non-physical realm. The weapons of cyber warriors – code – is physical at least in that it is made up of electrons, of physical matter with physical properties that obey the laws of physics. Likewise, the targets of cyber attacks, which can include programs, databases, websites, and so on, are equally made up of code or information also which are not adequately conceptualised as non-physical. Digital information is a particular arrangement of matter and any change to a program is a change to its underlying physical structure. Yet the claim about cyber attacks being non-physical is often elided with related but equally problematic claims, that they are non-kinetic or non-tangible for example.

[44] I am indebted here to Bernard Williams's discussion in his 'The human prejudice', in Adrian W. Moore (ed.), *Philosophy as a Humanistic Discipline* (Oxfordshire: Princeton University Press, 2005), pp. 135–52.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

Certainly it is true that the damage cyber attacks can cause is not due to kinetic energy, yet chemical and biological weapons do not cause harm via kinetic force either but their use is nevertheless still very much an aggressive act. And I can no more touch code than I can a laser or a high-energy electromagnetic pulse, yet the weaponised use of either may plausibly be conceived of as act of aggression.

Regardless, however, of whether we ought to consider cyber attacks to be physical or not, we go wrong if we think that this question has much bearing on whether a cyber attack can constitute a *casus belli*. Even if it were the case that cyber attacks are non-physical (or non-tangible or non-kinetic, etc.) this fact would actually be *completely irrelevant* to the question of whether they can represent aggressive acts or not. That cyber attacks are acts within the cyber realm is certainly a novel characteristic that distinguishes it from other attacks. But the domain in which those attacks take place are no more relevant to the question of whether they can constitute a *casus belli* as the fact that naval warfare takes place at sea or that an aerial campaign takes place in the sky (or, indeed, that attacks in one domain might have consequences in another). And we can quite naturally intuit this if we consider that we would be highly unlikely to assess the destruction of a nuclear power plant any differently if it was due to a cyber weapon or an actual bomb. Even if we granted that the former is non-physical this seems to have little bearing on our consideration that it would indeed, and just as with the use of a physical bomb, represent an act of war. The question of their physicality is simply an extraneous issue.

As for the issue that JWT is unable to accommodate the fact that cyber attacks target non-humans, this too seems both misguided and somewhat besides the point. The loss of human life might be sufficient, but it is not a necessary condition of a *casus belli*. Elaborating upon the above scenario, lets us further imagine that while the cyber attack causes irreparable damage to the plant without harming a single person, the conventional bomb both destroys the plant and causes numerous fatalities. There is no reason to think that the terms of just war theory are such that only the second scenario could constitute a *casus belli*. Insofar as the use of a cyber attack would presumably do indirect yet nevertheless significant damage to the interests of its citizens then it would still count as an aggressive act (a point we shall come on to shortly). Or consider an invasion by one state of an area of a neighbouring state that is devoid of occupants and completely lacking in cultural, financial, political, or defensive value for its people. It would be hard to say that a single human being is harmed in such an invasion, yet we readily recognise that this would represent a violation of the state's sovereignty and hence legitimate an armed response. In short, there is nothing in JWT that insists that human bodies must be the direct target of an attack for it to justify retaliation.

This leaves the issue of cyber attacks being non-violent. Again there is reason to doubt that this is descriptively adequate. We know from Stuxnet that cyber attacks can potentially be physically destructive in ways that we traditionally associate with acts of aggression, and where any future cyber attacks reach that level of violence then we can straightforwardly say that such an act would represent a *casus belli*. The only real question here is not would such attacks constitute a just cause but whether we can reasonably assume that the future of cyber conflicts will see the proliferation and use of Stuxnet-like weapons designed to cause physical damage or continue to be dominated by the largely disruptive attacks that have characterised most of the cyber attacks that we have so far encountered.[45] Either way, the relevant point for us is that cyber attacks can be violent in ways

[45] See May, 'The nature of war and the idea of "cyberwar"'.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

equivalent to the sort of violence traditional just war theory has hitherto been concerned with and, where that is the case, then we are able to assess such acts according to the same standards of aggression.

However, while it is possible to ground the idea of cyber *casus belli* in the same notions that ground our broader and more traditional accounts of *casus belli*, it nevertheless remains the case that many of the cyber attacks that we have experienced to date have been non-violent insofar as they lead neither to physical destruction nor human injuries or fatalities. They were purely disruptive. Could such non-violent attacks still constitute a *casus belli*? The answer to this is that they can, though we require a better sense of how non-violent cyber attacks can still be harmful in morally significant ways such that they are acts of aggression. And to recognise this we do need to go beyond traditional just war theory and accept that aggressive acts can take non-violent forms.

The way to get a handle on this is to consider both the nature of the computer systems cyber attacks target and the possible ramifications of those attacks. Remember that one problem we identified with an informational approach to cyber attacks is that it did not allow us to discriminate between the value or moral significance of different computer systems. An attack on *any* computer system seemed to legitimate a violent response. Clearly a somewhat more specific account of which targets are salient in this regard is required. What I want to suggest is that an attack can only rise to the level of a *casus belli* if it targets a state's computer systems necessary for the functioning of its *critical infrastructure*. This is not to deny that other computer systems can be the victim of a cyber attack; those of companies, universities, individuals, and so on, are regularly targeted. Rather, for the purposes of thinking about the legitimacy of an aggressive response, only attacks on the computer systems of critical infrastructure are potentially significant enough to warrant such a reaction. How can we justify this focus?

Definitions of critical infrastructure vary across countries but are united by the thought that the relevant asset must be 'vital' in order to count as critical. So the US Department of Homeland Security states that 'Critical infrastructure are the assets, systems, and networks, whether physical or virtual, so vital to the United States that their incapacitation or destruction would have a debilitating effect on security, national economic security, national public health or safety, or any combination thereof'.[46] Likewise, the UK government defines critical national infrastructure as: 'Those infrastructure assets (physical or electronic) that are vital to the continued delivery and integrity of the essential services upon which the UK relies, the loss or compromise of which would lead to severe economic or social consequences or to loss of life.'[47] While, therefore, any infrastructure sector (for example, communications, emergency services, energy, finance, food, government, health, transport, water) will be made up of numerous assets, it is only those the loss of which would have 'severe' or 'debilitating' consequences in terms of economic, social, or political disruption or the loss of human life that count as critical. This narrower definition is to be preferred to more expansive understanding of critical infrastructure as anything essential to the functioning of a society and economy, which then may, as a Chatham House report recently pointed out, raise questions as to whether the criticality of companies such as Google and Amazon to the functioning of a modern economy would also qualify them as critical infrastructure.[48]

[46] {http://www.dhs.gov/what-critical-infrastructure}.
[47] {https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/62504/strategic-framework.pdf}.
[48] Paul Cornish et al., 'Cyber Security and the UK's Critical National Infrastructure', *The Royal Institute of International Affairs* (Great Britain: Chatham House, 2011).

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

Both the US and UK definitions acknowledge that there can be 'virtual' or 'electronic' critical infrastructure assets, the damage or disruption of which would lead to 'severe' or 'debilitating' consequences. This infrastructure is considered critical because it is vital to the delivery of services necessary for securing and maintaining significant human interests in health, communication, energy, defence, order, and so on. Clearly only certain assets within a national infrastructure sector will be deemed critical depending on their function in providing these essential services and the potential consequences of their being disrupted. This allows for a certain amount of flexibility of judgement as to what counts as a critical infrastructure asset or not. It is also true that governments demarcate different infrastructure sectors (the US lists 16, the UK 9) but this reflects more the differing governmental structures of each country and the question of how departmental responsibility for each sector is appropriately allocated. It also speaks to the peculiarities of each country's needs and vulnerabilities. The US, for instance, lists the dams sector as critical infrastructure because of how reliant they are on dams for the provision of hydro-electric power, river navigation, water supply flood control, and waste management. The UK does not share such reliance and hence dams are listed only as a subsector of their water national infrastructure. But despite these variations, the point remains that critical infrastructure is identified with reference to their vital function in serving and realising significant human interests.

The focus on critical infrastructure and the disruption that cyber attacks can cause gives us a way of thinking in a quite straightforward and direct way about the potential harm those attacks might do. The 'wrong received' by a cyber attack lies in how disrupting the functionality of critical infrastructure has the potential to then (directly and/or indirectly) cause harm to the vital interests of human beings. And it is from this wrong that we are then able to make a judgement as to whether an attack constitutes a *casus belli* or not. Crucially, as opposed to the informational approach of IE, all this account needs in order to be both morally plausible and theoretically sufficient is a general notion of which human interests have adequate moral significance alongside an empirical understanding of how critical infrastructure assets relate to those in particular contexts. Both of these features are readily available within an anthropocentric LoA. The former underpins many contemporary 'interest-based' theories of human rights that explicitly do not seek to ground those rights in anything extra-human such as reason or natural rights,[49] while the latter is a matter only of identifying the network of institutions and relations between particular assets and particular human interests.

Cyber attacks can take both violent and non-violent forms. Where they are violent JWT can assess whether they constitute a *casus belli* in exactly the same way as it would a more traditional attack. In this sense there is continuity between the typical grounds for considering an act of aggression a *casus belli* in JWT and what might constitute a cyber *casus belli*. A slightly different account is required for non-violent cyber attacks. If non-violent cyber attack reaches a significant enough level of disruption such that they harm vital human interests then at that point they may be considered a *casus belli*.

A few points deserve clarification or further development. First of all, it is important to note that this account leaves no moral remainder; there are no aspects of the wrong of a cyber attack that remains unexplained, unjustified, or mysterious to us. For example, a further question that has vexed some commentators, is whether a cyber attack can be considered aggressive when the damage that it causes to a computer system may be temporary or completely reversible. We can now see why this is a misplaced worry. That a computer system itself may return to full functionality following an attack

[49] See, for instance, Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986).

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

proves to be of little consequence if what matters for the purpose of assessing whether it constitutes a *casus belli* or not is how, if at all, it harmed human interests. It may well be that a cyber attack permanently degrades the functionality of a computer system, but that is only relevant insofar as it means that the harm caused to human beings' interests is likely to be more profound than a brief attack where full functionality is quickly restored. Yet that need not always be true and we can imagine scenarios in which a brief or intermittent cyber attack may even prove more harmful than one that permanently disables a computer system. And so we go astray if we think that the question of the (im)permanence of the degradation to computer system functionality is central to how we conceptualise the aggressive nature of cyber attacks.[50]

It is important to recognise that while non-violent cyber attacks are aggressive insofar as they disrupt the functioning of critical infrastructure that are vital to securing and maintaining significant human interests, whether a cyber attacks represent a *casus belli* is a *judgement* that can only be made in relation to the consequences, potential and actual, of the attack on those interests. The conclusion of this argument is that cyber attacks (violent and non-violent) *can* represent a just cause, not that they necessarily *do*. The actual disruptive capacity of cyber attacks even on critical infrastructure is going to vary wildly. Some may cause very little disruption at all, either in terms of the duration of the attack or how it undermines the capacity of that infrastructure to function adequately, and hence be best thought of as a nuisance. Where that is the case, as may be true with most DDoS attacks, we are less likely to judge that such attacks constitute a *casus belli*. So there is good reason to think, as others have suggested,[51] that while Russia's cyber attacks on Estonia did not reach the level of a *casus belli* their attacks on Georgia a year later did, for instance. And we can make this judgement by comparing both the targets of the attacks and the severity of the disruption caused. If it turns out that in the future not many cyber attacks will ever meet the relevant threshold for being a non-violent *casus belli* then that is not an unfortunate outcome.

This raises the question of whether attacks on all critical infrastructure can potentially be considered a *casus belli*, especially as what counts as critical infrastructure does not only cover those assets that directly relate to the human interest in life and health. Alongside the Emergency Services, Defence Industrial Base, Healthcare and Public Health, and Food and Agriculture sectors, the US government also lists the Financial Services, Information Technology, and Transportation Systems sectors. Each of these clearly has an important role to play in underpinning those sectors that more directly address the basic interests of human life; without a functioning banking system people would not have the money to buy food, without a working transport system the emergency services would cease to function, and so on. Yet for each of those sectors mentioned, the government's justification for their

[50] This helps shed some light on exactly where advocates of the radical school go fundamentally wrong in their analysis. Firstly, they wrongly assume that cyber attacks cannot be violent in ways we would familiarly recognise as such. Stuxnet should be evidence enough that this is false. But, more importantly, they are simply wrong in thinking that because a cyber attack targets computer systems that we must account for the relevant harm in terms of damage to those systems, and hence need to create a new ethical framework in which we can conceptualise harm to non-humans, rather than to the human beings whose interests they are created to serve. If we avoid making that mistake then we can see how adopting an information perspective, either as an initial ethical starting point to which an anthropocentric LoA is added or as in some sense supplementing that human perspective, adds nothing to our understanding of the wrong of a cyber attack. Focusing on the (potential) consequences to human interests of disruptions to particular critical infrastructure assets is sufficient and provides all the reason we need to resist IE's argument that there are 'no good reasons to stop anywhere but at the bottom'. We can stop where we need to stop.

[51] See, for example, Lucas, *Ethics and Cyber Warfare*.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

inclusion as a critical infrastructure also refers to their vital social and economic functions. If we go back to the definitions of critical infrastructure offered earlier, it is important that the US includes 'national economic security' as an area of national life that critical infrastructure protects and services. And it is listed on equal terms alongside security and national public health or safety. The UK definition is even broader. It defines critical infrastructure in relation to those assets the disruption of which would cause severe economic, social, or political disruption, again alongside the loss of human life but without giving any sense of priority among them.

There are two issues here: can even severe economic, social, or political disruption constitute a *casus belli*? And could a cyber attack ever cause such severe disturbance? On the first question, there is obviously significant ambiguity as to what aspects of our economic, social, or political lives are pertinent here. While some people may be dependent (quite literally) on Facebook for much of their social interactions, we would not want to include any disruption to its services, no matter how severe, as a just cause for going to war (and that a Chatham House report even entertained this idea in relation to Google or Amazon is quite extraordinary).[52] Again, few commentators have agreed with the Estonian prime minister that the cyber attacks they were victims of during their conflict with Russia amount to a just cause, even though they disrupted the ability of the government to communicate with its people. Likewise Senator John McCain's declaration that Russia's alleged interference in the 2016 US presidential election amounted to an 'act of war' garnered little support, even though, if true, it may well be the case that Russia's involvement influenced the outcome.[53] Matters might be a little clearer in regards to economic disruption, but maybe only at the very extreme. If a cyber attack had the same catastrophic effects on human life and misery as Walzer sets out in relation to economic blockades, then we may follow him in thinking that it does represent an aggressive act.[54] But then imagine a cyber attack that, for instance, targeted the financial system and caused significant enough disruption, say through preventing electronic financial transactions, deleting bank account details, or corrupting information of capital holdings, to cost a national economy several billions. Would such an attack be deemed a just cause despite the fact that it does not cause physical harm? The answer to this is far from clear, though there is a *prima facie* case for thinking that it would insofar as such severe economic disruption would then have knock-on effects on almost all of the state's national interests and those of its individual members. Whether such an attack is within the realms of plausibility is, however, a matter of some considerable dispute even among technical experts.[55] So it is not quite clear what sort of credible economic, social, or political disruption could qualify as a *casus belli*, though this may represent more of a failure of imagination as to the extent of the disruption a cyber attack can cause than be an indication of the inherent limitations of cyber weapons.

All of this admittedly adds a somewhat consequentialist flavour to just war theory as applied to cyber attacks that may look incongruent with how it has been traditionally understood. Yet it would probably be a mistake to think that just war theory gives no weight to consequentialist reasons at all. It clearly must allow us to discriminate between, for example, the shooting of a bullet across a border and the launching of a nuclear missile, and it would seem that it is the varying likely consequences of

[52] Cornish et al., 'Cyber Security and the UK's Critical National Infrastructure'.
[53] {http://edition.cnn.com/2016/12/30/politics/mccain-cyber-hearing/index.html}.
[54] Walzer, *Just and Unjust Wars*, ch. 10.
[55] Gartzke, 'The myth of cyberwar'; Jon R. Lindsay, 'Stuxnet and the limits of cyber warfare', *Security Studies*, 22:3 (2013), pp. 365–404; Thomas Rid and Peter McBurney, 'Cyber-weapons', *The RUSI Journal*, 157:1 (2012), pp. 6–13.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

these attacks that does the work in justifying our intuition that the former *may* constitute a *casus belli* (depending on who fired the bullet at what, whether anyone was hurt, and so on) whereas the latter does without question. And anyway there are good reasons for thinking that such consequentialist reasoning should feature prominently in relation to cyber attacks, and that is because we know that many states have been hacking and attempting to hack into each other's critical infrastructure for many years now (what Lucas calls 'state-sponsored hacktivism').[56] Some of these have been discovered; it is reasonable to assume that there are many intrusions that have not been caught and which are ongoing. Even when such intrusions have been exposed it is very difficult to ascertain exactly what the perpetrators did while they had access to that system, whether they planted numerous 'logic bombs' to be 'detonated' at some strategic moment in the future, whether they exported critical data or altered significant information, or whether they were just relatively innocently poking around the system. It would therefore be precipitous to think that the mere intrusion into a critical infrastructure's computer system represented a just cause. Rather, we should view such activity as forms of espionage (which is not deemed to constitute a *casus belli*) until the point at which it is apparent that the intrusion was part of an attack that either causes severe disruption to that infrastructure or to make such disruption possible.

On a related point, the contemporary global situation today is characterised by a sort of ongoing but low-level cyber war, or maybe more accurately, a series of low-level cyber skirmishes between states. Some of these have been made public by governments themselves, like the Syrian Electronic Army's recent hack into the US army website or China's (alleged) intrusion into US federal government computers.[57] Others, such as the (again alleged) cyber attack by the UK spy agency GCHQ on a Belgian telecom firm, have only come to light through leaks such as the 'Snowden Files'.[58] Nevertheless, we can be certain that there is a lot more activity going on that we remain unaware of. War is such a terrible condition that it must be the aim of any ethical theory or legal framework to try to limit the circumstances in which it is justified to engage in conflict. If any cyber attack on a nation's critical infrastructure constituted a *casus belli* then the fact of the matter is that the international sphere would not be characterised by a Hobbesian state of nature but by a much worse situation in which the most powerful states would be facing each other all with right on their side. This would not limit warfare but provide a sort of universal just cause that all powerful states could claim as justification to go to war. The only way to avoid this unhappy picture is to think that the consequences of a cyber attack must have some role to play in our moral judgements of them.

Conclusion

For all the hyperbole about how the distinctiveness of cyber attacks and conflicts in the cyber realm leaves just war theory 'in tatters', this article has argued that, at least in the case of *casus belli*, the principles of just war theory can fairly straightforwardly be applied to the use of cyber weapons. The lesson we learnt from the failures of IE to provide a sufficient ethical framework for thinking about cyber attacks was that it is a mistake to think that we need to abandon the 'human perspective' in order to adopt an ontological framework somehow better suited to the peculiarities of cyber conflict. Indeed, we found that those ontological features of cyber attacks that are often used as justification for why JWT is an outdated framework in need of an upgrade – that they are non-physical, target

[56] Lucas, *Ethics and Cyber Warfare*.
[57] {http://www.bbc.co.uk/news/world-us-canada-33058755}.
[58] {http://www.bbc.co.uk/news/uk-24387578}.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press

non-humans, and are non-violent – were either descriptively dubious claims to begin with or, more importantly, simply not pertinent. It is true that many – maybe most – cyber attacks will be non-violent. But we should accept the possibility that even a non-violent cyber attack can rise to the level of a *casus belli* if it is disruptive enough that it significantly harms vital human interests. This, admittedly, is an amendment to how the just cause principle in traditional just war theory has traditionally been interpreted. And we must also admit the greater weight of more consequentialist reasons than the traditional account allows. But, while important, I take these to be amendments that remain within traditional just war theory rather than require its abandonment and the refounding of a new ethics of war more suited to the digital age. Indeed the broader point of this article is that while the security environment in which we now find ourselves is radically different we ought nevertheless to reaffirm traditional just war theory, as paradoxical as that admittedly sounds.

Acknowledgements

This research was undertaken during a period of research leave made possible by a British Academy Mid-Career Fellowship. The author would like to thank the British Academy for their generous funding and support.

Biographical information

Matt Sleat is a Reader in Political Theory at the University of Sheffield.

https://doi.org/10.1017/S026021051700047X Published online by Cambridge University Press