

Cyberspace and International Order

Madeline Carr

Note: this is pre-publication. Please do not cite.

Introduction

When *The Anarchical Society* was published in 1977, the world was on the doorstep of seismic technological change. Telephones were still attached to a cord, letters were a mainstay of private and commercial communication, and computers were housed in universities and military research facilities. Hedley Bull's ideas about how social processes between actors mitigate the anarchy of international politics were developed within the context of industrial age technology and he speculated only briefly on how emerging information and communications technology (ICT) might impact on those social processes in the future. Forty years of extraordinary development in technology – both in terms of scope and scale, have raised many questions, but fewer answers, about how emerging technologies reinforce or contradict what we thought we understood about international relations.

In fact, for most of those 40 years, policy makers (but less so IR academics) have been attuned to the imperative of trying to make sense of these technological shifts for conceptions of global security, power and order. There are too many significant policy implications to enumerate here but protecting critical infrastructure from cyber security vulnerabilities has been one of the more enduring concerns. Connecting critical infrastructure like energy plants, financial institutions, and transport systems to a global computer network comes with many opportunities for increased efficiency and lower costs which developed states have raced to exploit. But the relatively insecure nature of the Internet also introduces a range of vulnerabilities for those systems to be penetrated, manipulated, and damaged. The potential for such interference to lead to large scale destruction and loss of life, coupled with speculation that state or non-state actors may pursue that potential in lieu of, or in addition to, conventional kinetic force, has led to a focus on cyber security in domestic and international politics. This has manifested in a number of ways, including the development of national cyber security strategies, the establishment of dedicated cyber security military units and doctrine, and in discussions on international cooperation on cyber security at head-of-state level. States are clearly concerned about the implications of digital technology for violence and conflict in international relations. They are preparing both to respond to it and, one must surely conclude, to utilise it in order to pursue their own national interests.

Considering the very broad implications of this technology, there have been surprisingly few attempts to employ international relations theory, concepts, and ideas for understanding

the landscape of cyber (in)security. Most of the existing work on this emerges from scholars working on military doctrine or strategic studies with a particular (and somewhat repetitive) emphasis on the writings of Clausewitz.¹ In fact, a useful starting point for trying to systematically think through continuity and change brought about by the information age can be to return to some of the other enduring IR thinkers to consider how their ideas may help. In doing so, of course, one may also observe ways in which those ideas, useful and enduring as they might be, may also come under challenge from novel circumstances. With that in mind, this essay draws on several of Bull's ideas on international order to look specifically at the problem of *attribution* in cyberspace – the persistent difficulty of tracing activities in cyberspace back to a conclusively identifiable actor.

This problem of attribution presents a unique problem for international relations because it removes an important element of the social structure that has been part of what creates order, as Bull understood it. Bull's ideas about the states system, international society, and, therefore, international order, all rest on the (previously sound) assumption that state actors are readily and accurately identifiable and that, barring some exceptions, their actions are attributable to them. This is integral to states' ability to engage in practices like diplomacy and international law. The problem of attribution, then, introduces a new dimension of anarchy – of *disorder* - to the social practices of international relations.

Although attribution may have implications for all five of Bull's institutions of international society, I focus here predominantly on international law – specifically the laws of armed conflict and international humanitarian law – because this is a particularly active site of state cooperation and contestation about cyber security.² Bull's ideas about the goals of minimising violence and promoting peace are helpful here because both of them are major themes in debates about international law in cyberspace. This essay makes several observations about the rather mixed implications of the problem of attribution for international order. First, not only does attribution render international law difficult to apply in cyberspace, it also means that the reasons why actors tend to adhere to international law are weakened, and that the transition from informal international norms to customary law may be much slower than otherwise expected. At the same time, the challenges of this additional dimension of anarchy have been a catalyst for considerable progress in negotiations over norms of responsible state behaviour – an outcome that Bull himself might have anticipated. Finally, the potential for anonymity means that some unexpected and uncertain avenues for non-violent resolution of political tensions are developing.

¹ Notable exceptions include Joseph Nye (2010), Johan Eriksson and Giampiero Giacomello (2006), Mary McEvoy Manjikian (2010), Nazli Choucri (2012), and, most recently, Daniel McCarthy (2015) and Lucas Kello (2013).

² In fact, great power relations, diplomacy, and war are all central to the debates about international law and cyberspace and the essay offers some observations about them, but only in relation to international law. Each one of them (and possibly the balance of power as well) could be the focus of further study.

The essay is organised as follows. The next section outlines what Bull had to say about emerging technology in order to put his views into context and to delineate where and how his ideas fit in with the arguments put forward in this chapter. Following that, the essay briefly touches on why conclusive attribution of cyber attacks can be problematic and how scholars have dealt with this to date. It then moves to two substantive sections of analysis. The first engages with Bull's ideas about minimising violence and maintaining peace as goals of international society in order to establish the relationship between violence and cyberspace. The second substantive section deals with states' recourse to international law. The conclusions here are that the problem of attribution generates both opportunities and challenges for international order. If we are to maximise the first and minimise the second, it will be important to fully understand how they intersect with (and sometimes challenge or contradict) conventional ideas and concepts about international relations that developed in the context of industrial age technology.

Bull on Technology

It is perhaps more surprising that, in *The Anarchical Society*, Bull delved into the implications of what was then very nascent technological change to the extent that he did, rather than that he did not develop this aspect more fully. He was, after all, explicit that his inquiry into order was confined to 'enduring issues of human political structure' rather than the 'substantive issues of world politics' at that time (1977, xiii). In addition, the development of information and communication technologies (ICTs) in the mid-70s, when he was writing, pre-dated any real speculation on what networking technology would mean for international relations. And finally, there was very little movement elsewhere in the discipline of IR with which Bull might engage on these issues.

Bull's thoughts on the potential for ICTs to have a transformative effect on international relations are contained in Chapter 11 on the decline of the states system. His starting point is to engage with the debate of the late 1960s about the unifying or fragmenting influence of 'electronic' communication and media. Bull cites scholars like Brzezinski (1970, 3) who were arguing that the world remained fragmented despite the 'shrinking of the globe' while others (like McLuhan 1962) envisaged the future as one more united and, consequently, more peaceful. Bull explains that he finds the fragmentation argument more compelling because closer contact can generate new tensions and because he anticipated the benefits of new technology would be most pronounced at a national or regional level rather than an international level (1977, 273-4).

This debate about whether ICT's will bring us closer together, thereby rendering us more tolerant of one another's perspectives, or whether that proximity will exacerbate our differences and heighten tensions between us, continues well into the second decade of the 21st century. However, it is no longer the driving question at the heart of debates about technology and international relations. Nor, really, is the somewhat over-simplified question

that Bull poses of whether emerging technology spells the decline of the state. There are many ways in which states are choosing, or being forced to accept, further compromises of sovereignty that may eventually combine to significantly reshape our conception of what a 'state' is. On the other hand, despite the considerable role of US-based transnational private organisations, and the many ways in which civil society has been empowered, the states system continues, 'for the time being' as Bull would qualify it, to be the key mechanism for governing cyberspace (Carr 2016b). Today, scholars are more cautious about attributing deterministic outcomes to technology that not only evolves and changes very rapidly, but which is also no longer regarded (as industrial age technology was, to a large degree, in the 20th century) as a force for change that is divorced from human agency (Carr 2016a, 17-32).

Bull's analysis was narrowly confined to 'communications' technology and premised upon the *benefits* that he thought might derive from this technological shift. Neither he nor most other IR scholars at that time had any conception of the threats that would later come to be perceived as woven through so many aspects of politics, civil society, commerce, and military practice. Essentially, Bull regarded ICTs as another 'awkward fact' for the view of world politics as simply relations between states. But this, he pointed out was only consistent with a long list of anomalies and irregularities that had previously arisen and failed to bring about the decline of the states system (1977, 274). He also acknowledged, however, that 'a time may come when the anomalies and irregularities are so glaring that an alternative theory, better able to take account of these realities, will come to dominate the field' (1977, 275).

After 40 years of extraordinary technological change, it is clear that Bull's question about the impact of emerging digital technology on the states system was really a question too broad to be explored through such a narrow aperture. Engaging with a question like his in 2017 requires first addressing a whole field of constitutive issues that arise from what we now understand to be the complex interplay of politics and digital technologies. However, while Bull's analysis of technology may not have been particularly useful, his ideas about the social nature of international relations can certainly help us to begin working through those many granular questions that it is necessary to address in order to build understanding about international relations in the information age. Indeed, the problem of attribution is one such granular question and Bull's work facilitates an approach that brings in the social dimension of what has generally been regarded as an explicitly technical problem.

The Problem of Attribution

It can be difficult (sometimes impossible) to conclusively attribute cyber activity using technical methods (Wheeler and Larsen 2003). The skill of the attacker, the sophistication of the target's security architecture and practices, and the time between detection and investigation of an attack all present challenges to attribution. The fact that attribution is neither always *possible* nor always *impossible* has generated disagreements in IR about its

significance. We can only speculate on whether, in the future, technological solutions will be found to completely eliminate the attribution problem, or whether advances in shielding identities and masking actions online will keep pace with detection and tracing capabilities – such that the problem persists. The current state of play is one in which, for sophisticated actors (both state and non-state), it remains possible to avoid detection and conclusive attribution. For highly skilled security investigators, it is often possible to trace attacks to regions, states, or even neighbourhoods but not usually to make substantiated claims about the actor's identity or the motivation or intention behind an attack – a point I return to in the discussion of the application of international law.

Much of the literature on the attribution problem focuses quite narrowly on its implications for deterrence. Rid and Buchanan point out that attribution is 'at the core of virtually all forms of coercion and deterrence'. They regard it as impacting on a state's 'credibility, its effectiveness, and ultimately its liberty and its security' (2015, 4). Clark and Landau write that '[a]ttribution is central to *deterrence*, the idea that one can dissuade attackers from acting through fear of some sort of retaliation. *Retaliation requires knowing with full certainty who the attackers are*' (2011, 25 italics in original). The challenges of deterrence in cyberspace coupled with the capacity for less conventionally powerful actors to exploit cyber vulnerabilities has caused concern in many quarters (while clearly being recognised as an opportunity by others).

In *The Anarchical Society*, Bull's arguments on deterrence, its role in balancing power, in rendering war irrational, and in preserving peace, all rests on the premise that states are identifiable (1977, 117-26). He considers two technological developments that might upset mutual (nuclear) deterrence: the acquisition of perfect defence and the capacity to disarm an opponent's retaliatory forces (1977, 124-5). In the context of cyber security, the first development is an ongoing pursuit in which operators of computer systems and networks maintain a regime of constantly updating and patching their systems and improving practices so as to minimise the likelihood of penetration. However, at this stage, there really is no expectation that *any* network is impervious to intrusion and exploitation. Bull's second development, of course, returns us to attribution – because we need to know *who our opponent is* in order to disarm them.

Deterrence is not always effective in the physical world. When dealing with security threats from non-state actors, for example, the kinds of coercive mechanisms that work on state actors have failed to change the behavior of those who feel they have nothing to lose or who actively seek martyrdom. In the context of cyber security, however, it is proving particularly frustrating for many states that they are unable to deter other *state* actors that are able either to hide their actions completely or to mask them behind proxies (by contracting private hackers). Frustration with the limitations of deterrence in cyberspace, where even great powers seem unable to have their way, has led to perhaps one of the more troubling developments in the literature on attribution. Some analysts have suggested

that evidentiary standards for the attribution of cyber attacks be reconsidered so as not to require conclusive technical proof (Healey 2011, Knake 2010) – another proposal that I return to in the discussion on international law.

This literature on attribution and deterrence is important because it points to the challenges of continuing to rely upon a mechanism that has been important throughout the industrial age (and before) but which, it appears, may have limited utility in the information age. It is also important because it is by far the dominant approach of social science scholars interested in the political implications of the attribution problem. At the same time, however, this persistent linkage of attribution to deterrence has tended to limit the parameters of the debate to strategic issues rather than the broader social and political factors that make attribution attractive or desirable in the first place. By continuing to focus so specifically on how attribution can or cannot be reconciled with deterrence, we risk missing the broader implications of anonymity in cyberspace for international relations.

Attribution and International Order

As pointed out in the introduction, this essay focuses on international law not because it is the only one of Bull's institutions for which the problem of attribution has implications, but because it is the site of much of the current international cooperation and contestation around cyber security. Since this analysis is concerned with the laws of armed conflict, and since the potential for cyber attacks to result in physical violence has been quite vigorously disputed, it is necessary to establish the connection between the potential for violence and cyber security. Bull's holistic approach to the goal of minimising violence – one that transcends but also incorporates international relations – helps us move away from strategic questions about attribution (How can we develop deterrence? How can we solve the attribution problem?) to think more clearly about what causes different actors to perceive the threat of violence in cyberspace differently. It also opens up space for considering other approaches to cyber capabilities including one that promotes peace by allowing for non-violent solutions to political tensions.

The Goal of Minimising Violence

Bull argues that all societies seek to a) ensure that 'life will be in some measure secure against violence resulting in death or bodily harm', b) 'ensure that promises, once made, will be kept, or that agreements, once undertaken, will be carried out' and c) ensure that the 'possession of things will remain stable to some degree, and will not be subject to challenges that are constant and without limit' (1977, 4-5). In Bull's view, these goals, 'life, truth and property', are *elementary* goals. Without them, he suggested, we could not call a group a 'society'. He argued that they are also *primary*, because all other goals that societies may have presuppose these ones, and that they are *universal* because all societies seem to 'take account of them' (1977, 5-6). The extent to which the elementary, primary, and universal goal of minimising violence is relevant to cyberspace is by no means settled in the

scholarship. Some regard the potential for large scale devastation and loss of life as very worrying while others feel that this is a remote and unlikely threat. While these debates are quite polarised (and somewhat stagnant), they both offer important observations and conclusions.

Violence matters – the ‘Cyber Pearl Harbor’ view

Since the mid-1990s, policy makers have been concerned about the prospect of a large scale and violent cyber attack on *critical infrastructure* – those systems like power, water, and communications that we regard as essential to the smooth functioning of society. They have consistently expressed concern that decades of privatisation of critical infrastructure combined with reliance on insecure networked systems is a dangerous development, pregnant with the potential for physical destruction and loss of life (Carr 2016b; Legrand 2014). In 2012, Leon Panetta spoke about the threats to US critical infrastructure as he perceived them. He explained that as ‘director of the CIA and now Secretary of Defense, I have understood that cyber attacks are every bit as real as the more well-known threats like terrorism, nuclear weapons proliferation and the turmoil that we see in the Middle East’ (2012). The 2014 NATO summit declaration stated that ‘[c]yber attacks can reach a threshold that threatens national and Euro-Atlantic prosperity, security, and stability. Their impact could be as harmful to modern societies as a conventional attack’ (NATO 2014).

National cyber security strategies tend to focus on the threat to critical infrastructure and the potential for interference in their command and control systems to have catastrophic effects – the release of water from a dam to flood a populated valley, for example, or the interruption of power supplies in a dangerously cold winter. These scenarios are sometimes referred to as a ‘Cyber Pearl Harbor’, reflecting political leaders’ anxieties about being taken by surprise by a devastating attack and being underprepared. Despite these strong views and persistent concerns, it must be noted that there are others who argue that digital technology is not able to deliver violence and is therefore much more limited in its utility and threat.

Cyber is not violent – the sceptic view

When arguments are made to the effect that the threat of violence from digital technologies is inflated, they are very often framed in terms of a belief in the continuity of the relationship between technology and global affairs – the status quo. This argument rests on the premise that despite the vulnerabilities to critical infrastructure, the reliance of civilian and military systems on the Internet and other networks, and the occasional intent of actors in international relations to cause physical destruction, loss of life, and/or large scale disruption, we have yet to experience a cyber attack with these effects. Therefore, some scholars argue, there is no basis upon which to expect that we will experience one in the future. Thomas Rid has written extensively (and sceptically) about violence and cyber attacks. He recognises that the kind of critical infrastructure attack about which others are so concerned is possible, but emphasises that, ‘so far, no such scenario has ever happened...

Not a single human being has ever been killed or hurt as a result of a code-triggered cyber attack' (2013, 13).

The status quo argument required clarification in the wake of the Stuxnet operation, revealed in 2010, in which a nuclear enrichment facility in Iran was damaged by a computer worm (Clayton 2010). Although this incident did not result in any loss of life, the implications of an actor penetrating a highly secure facility using cyber tools and causing physical destruction to critical infrastructure is the very scenario that keeps policy makers awake at night. There are a number of ways in which scholars sceptical about the destructive potential of digital technology have responded to this important example. They argue that: a) the attack was very expensive and consequently beyond the reach of, or unattractive to, most actors (Rid 2012; Lindsay 2013, 388); b) the attack was limited in its effects because it did not destroy the Iranian nuclear facility and therefore is unlikely to lead to more widespread use of similar exploits (Lindsay 2013, 390-2); and c) that cyber weapons like Stuxnet are 'use and lose' capabilities which must be deployed in secrecy and so have little to offer in terms of compellence or deterrence. (Gartzke 2013, 60).

There are weaknesses in all three of these assertions. Costs are not static, limited efficacy in one attack is no indication of future developments, and claims about the limited appeal of cyber weapons need to be substantiated by some kind of empirical research. However, if this line of argument – that digital technologies *are* limited in their potential to cause violent acts – were more robust, then we might assume that attribution matters less for sustaining the goal of limiting violence and, by extension, international order. Unfortunately, the flaws in these arguments aside, a determinist approach to technology has little utility in international relations because it fails to take into account human agency and it also discounts the extent to which international relations itself shapes technology (Carr 2016a).

Although cyber security fears (like any other fear) are no doubt over-inflated by some actors in some circumstances and for some purposes, that does not calm the nerves of policy makers who are alert to the potential for technological vulnerabilities to be exploited by those who might wish to pose a serious challenge to international order. These two dichotomous positions are unlikely to be reconciled in the near future. Sceptics would only be convinced by a proliferation of devastating attacks, while few policy makers are likely to feel easy about ignoring the threats just because they have not yet eventuated. There is a third possible approach that moves beyond both of these positions: that digital technologies can also provide much less violent solutions to political conflict.

The potential for maintaining peace

In addition to pointing out the possibilities of ICTs for physical destruction, the 2010 Stuxnet attack also precipitates consideration of the potential for digital technologies to be employed to address political tension *without* violence and this therefore has implications not only for the future of political conflict, but also for the maintenance of peace, Bull's third goal of international society. Although this is a less widely held view than either of the

above approaches to violence and ICTs, it is worthy of some consideration here and especially in light of Bull's point that he was referring not to a 'universal or permanent peace' but rather to the absence of war as a 'normal condition' (1977, 18). Once again, Stuxnet provides a useful example through which we can explore this but with an important caveat.

Despite the wide coverage of this event in academic literature and in the media, it is important to keep in mind that our knowledge of the Stuxnet attack remains largely anecdotal and unconfirmed (a key problem for scholarship in this field). There has been one dominant narrative to develop after Stuxnet became public knowledge and this is based on the work of an award winning American journalist, David Sanger, who has based his account on interviews with many high level (but attributable) sources. For many reasons, relying upon 'evidence' like this is deeply problematic but, for the purposes of this chapter, I do not engage with it as 'truth' but rather as a useful hypothetical. It does not matter for the purpose that it will be used here, whether Sanger's account is completely, partially, or not at all accurate. What his account offers this essay is a platform to think through possible implications of similar attacks.

According to Sanger, Israel's growing concerns about the Iranian nuclear program were edging the state toward plans for a kinetic attack on the Natanz nuclear facility. In an effort to prevent action that they felt may lead to a catastrophic conflict in the Middle East, Sanger suggests, the US worked with Israel to develop the 'Olympic Games' program, of which the Stuxnet worm was a central component. The program ran for several years, not destroying the nuclear facility but delaying progress sufficiently to slow down Iran's transition into a nuclear state. In this way, then, with no loss of life and no escalation to a kinetic conflict, the Stuxnet worm potentially delivered a non-violent solution to an extremely dangerous, volatile, and potentially devastating political crisis in the Middle East.

Although there has been widespread conjecture (supported by Sanger's story) that the US and Israel were behind Stuxnet, neither state has claimed responsibility. This ambiguity possibly left a wider range of response options open to Iran. If it were conclusively attributed, Stuxnet may have forced a different response not only from Iran and its supporters, but also from the rest of the international community. It is difficult to address the issue when there is no conclusive identification of the actors behind Stuxnet and, in that way, the problem of attribution may have provided a pressure valve for Israel and the US as well as for Iran and for the rest of the international community. It is possible that Stuxnet represented a very creative approach to promoting peace (or the absence of war) and that cyber capabilities have much more potential to do so than those who argue about their (non)violent properties acknowledge.

Bull's rationale for the goal of minimising violence is that unless people enjoy some 'measure of security against the threat of death or injury at the hands of others, they are not able to devote energy or attention enough to other objects to be able to accomplish

them' (1977, 5). This is reflected in comments by the Russian Foreign Minister, Igor Ivanov, to the UN in 1998. Ivanov wrote to the Secretary General to express Russia's concerns about the potential for ICTs to undermine international order. He wrote that, in Russia's view, the international community must not 'permit the emergence of a fundamentally new area of international confrontation' that would 'divert an enormous amount of resources that are so necessary for peaceful creativity and development' (in Tikk Ringas 2015). Tikk Ringas cites this as the first instance of a state actor linking ICTs to international law in the context of global security and suggests that Ivanov's letter was the genesis of significant discussion amongst the great powers about international law and cyber security (2015).

International law and cyberspace

There has been a debate about whether international law conceived of in a different technological age, could be readily applied to cyberspace. On the one hand, the International Court of Justice states that the laws of armed conflict apply to 'any use of force, regardless of the weapons employed', and on the other, the Permanent Court of International Justice states that acts not forbidden in international law are generally permitted (Schmitt 2013, 3). Although both Russia and China have argued strenuously for a treaty to address global cyber security concerns, the US view that no new law is necessary has thus far prevailed.³ The focus of international negotiation and discussion has, instead, revolved around two axes: first, establishing whether and how existing international law applies in cyberspace; and second, negotiating norms of responsible state behaviour. The expectation is that some or all of these may one day crystallise into customary law. As Bull noted in regard to the different problems of his time, the value of international law lies not in its capacity to dictate rules that states must adhere to and to stipulate consequences for the violation of those rules. Rather, the value of international law lies in its capacity to provide a mechanism or a channel through which agreed interests may be institutionalised, acknowledged, and organised. In doing so, international law provides some measure of predictability and reassurance about state behaviour. It allows states to signal their 'intentions with regard to the matter in question' (1977, 142).

When thinking through the implications of the problem of attribution for international law, there are three important questions that Bull's conception of this institution of international society raises. First, how can the law be applied to anonymous actors? Second, how does the problem of attribution impact on the motivation for actors to abide by the law? And finally, how useful is international law, as Bull conceived it, for signalling states' intentions and for promoting predictability? Before engaging with these questions, there are some definitional issues that produce real impediments to applying international law to cyberspace. Understanding these is essential to comprehending the complexity of these three questions.

³ One exception is the Council of Europe Convention on Cybercrime, but this does not address global security concerns.

Applying international law in cyberspace

Although states agreed in 2013 that existing international law *does* apply in cyberspace, the problem of *how* to apply it has yet to be resolved. Questions persist about the interpretation of Article 2(4) of the UN Charter, which prohibits states from the ‘use of force’ unless granted authorisation by the UN Security Council or unless (as stipulated under Article 51) responding to an ‘armed attack’. Exactly how to define an ‘armed attack’ and what exactly the threshold for ‘use of force’ should be has confounded legal scholars in the context of ICTs. These concepts, upon which the laws of armed conflict and international humanitarian law rest, were developed prior to the advent of modern ICTs and, consequently, they have proven exceptionally difficult to map onto the complex nature of cyber incidents. Boothby et al point out that an armed attack should be ‘grave in scale and effects’, though there is no test to distinguish ‘grave’ from ‘non-grave’ consequences (2012, 83). Most legal experts, they suggest, agree that an armed attack will result in ‘death or a significant degree of injury to persons or physical damage to property’ (2012, 83).

On the one hand, this returns us to the arguments of the cyber sceptics who will point to the fact that there has never been a cyber attack that resulted in death or injury on such a scale. But what, then, of the Stuxnet attack, of which former CIA director Michael Hayden has said ‘...you can’t help but describe it as an attack on critical infrastructure’ (cited in Farwell and Rohozinski 2012, 111)? Unlike kinetic weapons, cyber tools can be designed to cause large scale physical damage to critical infrastructure *without* killing people. That is a unique capability that one could argue avoids traditional interpretations of ‘armed attack’, but it is not necessarily a practice that states would regard as permissible. In arguing that some of our ideas about violence and war may require rethinking in the information age, Chris Demchak has called attention to what she calls ‘wars of disruption’, in which the focus is no longer *lethality* but organisational *disruption* through information systems (2011). Some certainly feel that there is a gap here between what was intended in *lex lata* and what current circumstances call for and this has been at the heart of these debates about definitions.

It is reasonable to expect that, some time in the future, these questions will be satisfactorily resolved – that we shall see some consensus on how to define the ‘use of force’ and ‘armed attack’ in cyberspace. However, as long as they are defined by *consequences* (ie loss of life, large scale disruption), the problem of attribution will still act as an impediment to the application of international law. Consequences themselves are not adequate because it is the actor and their motivation that combines with consequences to allow us to classify and make sense of any kind of violence. Without a clear identity of the perpetrator, it can be very difficult to separate criminal activity from politically motivated activity, or to separate politically motivated activity undertaken by a non-state actor from that of a state actor. And that in turn raises questions about what type of law applies and what kind of penalty is appropriate or legal.

Even if we put to one side the current challenges of interpreting international law in this context and focus again on the problem of attribution, a second question arises: what will motivate states to adhere to the law if they may violate it in anonymity? If it is possible to carry out illegal acts such as attacks on critical infrastructure without those actions being conclusively attributed and, therefore, without eliciting the usual consequences of violating international law, what might restrain states from exploiting the opportunities that ICTs present?

Why obey the law?

Bull articulates three reasons why states obey international law. First, because the law may be regarded by them as 'valuable, mandatory or obligatory', second because of the threat of coercion, and third in the hope that doing so may prompt reciprocal behaviour from other states (1977, 139-140). He draws these social factors back to the self-interest (or national interest) of states by pointing out that the 'importance of international law does not rest on the willingness of states to abide by its principles to the detriment of their interests, but in the fact that they so often judge it in their interests to conform to it' (1977, 140).

Certainly, political actors may abide by international law in cyberspace so as to signal that their state upholds its obligations, so as to avoid coercion, and in the hopes of fostering reciprocity. However, this conception of the social nature of international law is predicated on a clear understanding of who is acting and who is being acted upon and this is deeply problematic in an anonymous environment. States may very well have a shared interest in protecting critical infrastructure from cyber attacks, but if anonymity is an option, Bull's motivations for adhering to the law are no longer as compelling an explanation of state behaviour as they might otherwise be. For example, there has been speculation that the attacks on the Ukraine power grid in December 2015 were state-initiated or state-sponsored. If that were true, it may not indicate that the offending state had little interest in a prohibition on critical infrastructure attacks. It may instead indicate that although they shared that interest, there was some expectation that by avoiding conclusive attribution, they could also satisfy other foreign policy interests without cost. The problem of attribution complicates these notions of obligation, coercion, and reciprocity that Bull sees as fundamental to the motivation of states to adhere to the law.

International law in the information age

Given the above challenges that attribution introduces to *applying* international law and given the ways in which it may reduce actors' motivations for *adhering* to the law, we may still ask the following question: how effective is international law as a mechanism to promote predictability and reassurance about state behaviour? At this point, it is useful to consider some of the progress that has been made around agreeing on cyber norms.

In response to Ivanov's 1998 letter to the UN, the UN Disarmament Committee established the *UN Group of Governmental Experts on Developments in the Field of Information and Telecommunications in the Context of International Security* (UNGGE). Since 2004, this has

been the primary site of global political debate about international law and cyber war and it has resulted in progress on establishing some norms of responsible state behavior in cyberspace. The 2015 UNGGE meeting produced a consensus report that acknowledged that ‘the use of ICTs in future conflicts between States is becoming more likely’ (United Nations 2015, 8). The report proposed 11 voluntary, non-binding norms, rules, or principles of responsible behavior for all states (including attacks on critical infrastructure) (United Nations 2015, 8). These norms, the report explained, are aimed at ‘promoting an open, secure, stable, accessible and peaceful ICT environment’ (United Nations 2015, 6).

It is not always the case that norms make the transition to customary law (and nor is it necessary for them to do so in order to be effective in shaping state behaviour or expressing shared interests) (Erskine and Carr 2016). However, within and around the UNGGE process there is some expectation that these norms of responsible state behaviour in cyberspace will become sufficiently embedded in, and representative of, state practice that they will eventually be recognised as customary law. Bull is very clear that in assessing the efficacy of international law it is not necessary to find that states always adhere to it and never violate it. Indeed, he makes the observation that ‘in cases where conformity between actual and prescribed behaviour can be regarded as a forgone conclusion, there can be no point in having rules at all’ (1977, 136). Rather than evaluating the extent to which actors’ behaviour is shaped by laws, Bull suggests that the question should be whether international law is observed to a sufficient degree to be regarded as a means of preserving international order (1977, 137). Here, the problem of attribution raises a unique problem in that it can be difficult to determine whether states are, indeed, exhibiting some generality of practice which might then be considered to be indicative of customary law or whether they are *claiming* to do so while in fact regularly violating those norms without being detected. This, Boothby suggests, could mean that customary law in this context is very slow to develop and we face an extended period of uncertainty about how effectively international laws of cyberspace contribute to international order (2016).

What this UNGGE process *has* done, however, is provide a mechanism for states, especially the great powers, to express their views, articulate their interests, and negotiate both the common ground upon which they agree and also those divergences that are so fundamental as to prevent further progress on the questions discussed above. One of these divergences is certainly approaches to attribution and this comes through clearly in the 2015 UNGGE report.

In addition to articulating state concerns that ‘the misuse of ICTs may harm international peace and security’ (2015, 6), the report also makes two important statements that reveal the great power tension around the problem of attribution. One of the proposed norms reflects the (largely Western) view that technical attribution is too difficult and uncertain to be considered essential to retaliation. The report therefore proposes that, ‘[i]n case of ICT incidents, States should consider all relevant information, including the larger context of the

event, the challenges of attribution in the ICT environment and the nature and extent of the consequences' (2015, 7).

This view reflects the frustration of some states that have been unable to deal effectively with ongoing cyber attacks that they believe are state sponsored or state supported. Neither technology nor the law has proved adequate for protecting state assets – a particular frustration for the US. Jason Healey has argued for avoidance of the trap of 'attribution fixation', by which he means 'the belief that [analysts] cannot assess which organization or nation was behind an attack until technical forensics discovers the identity of the attacking machines'. He suggests that 'attribution becomes far more tractable when approached as a top-down policy issue with nations held responsible for major attacks originating from their territory or conducted by their citizens' (2011, 1).

Essentially, this approach to attribution is one that avoids the difficulties of the accurate forensic analysis of cyber incidents through technical means and relies instead upon judgements about who one feels was *most likely* behind the attack, given a whole range of other factors like capability and motivation. The momentum behind this position in the US and its expression in the UNGGE norm signal the somewhat concerning potential for states to talk themselves out of the necessity of the burden of proof, or of establishing lower evidentiary standards for attributing cyber attacks, thereby opening up the way for unsubstantiated accusations, allegations, and even misdirection of blame for malicious actions.

Recognition of these dangers is noted later in the same report, reflecting the views of Russia and China. In the discussion on how international law applies to the use of ICTs, the report states that '...the indication that an ICT activity was launched or otherwise originates from the territory or the ICT infrastructure of a State may be insufficient in itself to attribute the activity to that State. The Group noted that the accusations of organizing and implementing wrongful acts brought against States should be substantiated' (United Nations 2015, 13). This is an explicit response to US allegations of illicit behaviour on the part of the Chinese and Russians that have not been accompanied by convincing attributory evidence. Examples of these include the 2014 indictment of five serving Chinese military officers over charges that they had been responsible for a sustained campaign of Chinese industrial espionage (Department of Justice 2014) and the 2015 imposition of (additional) sanctions against North Korean officials in response to the alleged attacks on Sony Pictures (Obama 2015c). Some regard these US responses as an effort to send a message that they will not continue to tolerate violations of their sovereign cyberspace.

This tension between the great powers over attribution is one of the major fissures that prevents more forward momentum on the applicability of international law in cyberspace. It also raises questions about the efficacy of international law, not as a means of enforcing rules (it is clear that attribution is necessary for that), but, as Bull suggested, as a mechanism for institutionalising shared interests, for signalling state intentions, and for promoting

predictability. This may slow down or arrest the development of further much-needed clarity in customary law.

Conclusion

There are a number of conclusions to be drawn from this analysis, both for how we think about the problem of attribution in international relations and for how Hedley Bull's ideas about international order stand up in the information age. His work is particularly useful for discussing attribution because it helps us to consider it and to analyse it as a *social* dimension. Engaging with his ideas about international order opens up a whole new landscape for thinking about an issue that has previously been considered almost exclusively through a technical or a strategic lens. Recognition of actors and the presumption that states are clearly identifiable is central to so many aspects of international relations that the inability to do so adds as a new dimension to the anarchy that (to some extent) shapes global politics. Perhaps most significantly, Bull helps to redirect the focus from looking for solutions to looking for the right questions.

In addressing the long-standing question of whether or not cyber attacks can be violent, Bull's work allows us to break out of the rigid confines of the dominant debate about whether cyber tools can result in violence or not. Instead, by taking into account his views on peace, we might look instead at the potential for ICTs to be employed in resolving political tensions in cyberspace and the potential for unattributed actions to further minimise confrontation.

In terms of international law, the implications of attribution are mixed. Fundamentally, it makes the applicability and the motivation for actors to comply deeply problematic. It also means that it will be very difficult for us to *recognise* customary law if and when it does develop out of the proposed UNGGE norms. On the other hand, confronted with these challenges, with this new dimension of anarchy, we see states responding very much as they have in the past, by balancing opportunities to exploit the potential of new technologies to pursue their national interest with participating in social practices like negotiation, diplomacy, cooperation and, however slowly, agreements on responsible state behaviour in cyberspace. And this would possibly be more or less as Bull might have expected.

Finally, having weighed up these factors of continuity and change, one would have to say that ICTs could no longer be regarded as simply another 'awkward fact' for international relations. Rather, if Bull were working now, he might come to regard them as a central 'issue of human political structure' (1977, xiii). They are an integral element that is woven inextricably through many aspects of our civil, political, military, and commercial existence and they have as much – but quite possibly more – potential to impact upon international order as industrial age technology did. It is impossible (and unwise) to make predictions about how this technology will be deployed in the future, and about what the long term implications for international order will be, and I think that if Bull were writing now, he

would widen his lens to explore the intersection of technology and international relations much more carefully. Instead, it falls to us to conduct further work that engages with ideas, concepts, and methodologies from international relations and other fields in order to better understand the massive technological shift that we are now living through.