Cyber-weapons are called viruses for a reason: Statecraft and security in the digital age

Inaugural lecture by Ciaran Martin, King's College London November 2020

Can I first thank King's College for two things? The first is the honour of a visiting Professorship at such a prestigious institution. The second is the privilege of being able to talk to you tonight, albeit virtually.

My appointment as a visiting Professor is a joint initiative between two distinctive parts of the University, both of whom enjoy global reputations in their own right. One is the War Studies Department, which counts, amongst its many other distinguished luminaries, someone I consider to be a long-time mentor in both Government and academic work, Sir David Omand. It is daunting for me to follow in his footsteps and deliver a keynote lecture here. The other is the Strand Group, one of the pre-eminent institutions for the applied study of British Government, with a veritable Hall of Fame of former Whitehall mandarins and advisers. I am exceptionally grateful to Dr. Jon Davis for his encouragement and for bringing me into the group.

The combined appointment is significant for me for other reasons.

I believe a secure technological environment is vital for our future prosperity. This horrible pandemic year of 2020 is the ultimate proof of that. Technology has kept us at least partly going professionally, and has kept us in touch with our loved ones.

Think of how this pandemic would have been 30 years ago. Or think about what it would have been like this year if technology couldn't have coped with the increased demand, or if repeated security disasters had destroyed people's confidence in using it. A terrible year would have been a whole lot worse.

So we must maintain trust in cyberspace. And I believe stewardship of the digital environment straddles both national security and those parts of Government and wider society that deal with the mundane realities of everyday life.

That's why we built the UK's National Cyber Security Centre in the way we did. Yes: it can covertly detect and deal with the most potent states. But it also stops fake emails from criminals pretending to be the tax or social security department, and it tells people how to work from home safely.

Here's an example of that dual purpose in action from an intense but not untypical day just before lockdown in March. I started my working day with a classified discussion of a nation state intrusion. For obvious reasons I am not going to say anything about the detail. All I will say is that there was the usual discussion about how we detected it, what the risks to the UK were, the mitigations, the case for public attribution (the efficacy of that is a discussion for another day) and, importantly for this discussion, the consideration of response options and discussions of our own cyber capabilities. In other words, the classic playbook. There were the usual references to the important contemporary concepts of 'full spectrum response' and 'the grey zone' and so on.

I had to leave the a meeting a little early. This was because I was booked to go to a BBC studio to sit on the sofa of the brilliant Victoria Derbyshire on her live show. I was there to explain to her viewers how to protect themselves and their families following the discovery of some technical flaws in the hardware of internet-enabled cameras on home devices. The flaws, though not widespread, were real. There'd been a handful of upsetting cases in the US where hackers had harassed vulnerable children through the cameras that concerned parents had placed in their in their bedrooms (something that is increasingly common in respect of older children with physical disabilities or mental health conditions; children who might be at risk of self-harm, or falls, for example).

So I was there to provide basic advice on how to reduce the risk of falling victim to this sort of attack. I was also there to reassure British citizens that (a) most cameras were perfectly safe and (b) that the Government planned to legislate for higher standards in such Internet of Things (or IoT) devices to make sure in the future it would be much harder for this minority of faulty products to get onto the market.

The contrast between the two morning tasks could not have been greater. But it was emphatically all part of the same job. Detection of covert nation state intrusion and the daytime TV sofa is cyber security in microcosm; the full spectrum of the challenge of cyber for Governments in a couple of hours.

At one end of the spectrum, the first meeting illustrated that cyber is now a strategic domain of operations for nation states. The digital domain is now one where nation states seek to project power and seek strategic advantage.

That involves things like large scale espionage, influencing operations, including in political discourse, intruding onto networks to 'pre-position', in the jargon, for a possible future disruptive attack. Occasionally this aspect of cyber statecraft involves actual disruption, sometimes unintentionally. It could (and there is increasing discussion about this) be a major aspect of future conflict.

For all these reasons, most major nation states think about the cyber domain as a domain of operations. Indeed NATO formally designated cyber as a domain in operations in 2016. That must be right.

But the second appointment, a message of basic consumer protection, reflects, in my view, the ***primary*** characteristic of cyber as a domain.

Of course the Internet, and other new technologies, have military and national security roots in part. But the modern digital domain is a place of social interaction, information exchange, debate, and very, very large scale commerce. Whatever the legitimate concerns about online harms, it remains, overwhelmingly, a domain of peaceful social and economic activity.

It is a ubiquitous, constant experience for us all. As and when you get bored during this lecture and get out your hand-held device to text a spouse, check the latest from the US counts and courts, buy something, enrol for the next King's event, or tweet how much you disagree with this lecture, turn the heating on (for the very wired-up amongst you) or whatever you may be doing, you are a civilian in the cyber domain even as you ponder its strategic national security implications in the context of this talk.

The two sides of technological security have the same core technological principles. We are talking about, for the most part, the same technology. The concepts of resilience, vulnerabilities, exploits, disruption, exfiltration of data apply throughout.

Cyber as a domain of military and national security operations co-exists with cyber as a domain of everyday life. It's the same domain.

And I think the experience of the NCSC shows that when it comes to thinking about *defending* the cyber domain, we get this point, for the UK at least.

However, as we in the West move more towards thinking of cyber as a domain of operations, there is more and more talk of the *projection of our own cyber capabilities* in the cyber domain.

This is a necessary part of modern statecraft. But as we develop our thinking, the same risk arises again; the risk of failing to realise the fundamental point that the domain of operations and the domain of peaceful activity are inseparable.

So, frankly, I worry that we are having two by and large largely separate conversations.

One conversation is amongst a national security and defence community. It is about cyber as a domain of operations.

Another is among civilian technologists. It is about securing the digital environment.

Both conversations are necessary, and both are entirely legitimate.

But they are separate. They are mostly unconnected, save for sporadic incidences of the two groups talking past, and occasionally shouting at, each other. They should be two parts of the same conversation about the same technologies and the same interaction with that technology of the same human beings.

So I hope that this event, jointly hosted as it is by War Studies and the Strand Group, and my other work here and in my main job in Oxford, plays a small part in trying to bring the two concepts together into an integrated consideration of statecraft in the digital age.

And I want, in the rest of this inaugural lecture, to look at the ways in which Western states are seeking project themselves in the cyber domain in the national interest. I want to look at the efficacy of the various approaches, and the risks. And I want to consider how this should fit with our overall objectives for how we want to act as free, open, liberal and democratic nations in the cyber domain and how relates to our own citizens' security in cyberspace.

***

To help the discussion, I want to introduce, or arguably, reintroduce, two concepts.

The first is cyber not just as a domain, but as an environment. It is so ubiquitous in our everyday life there is a strong case for this type of analogy. Cyber is a domain. But it is more than a domain.

If the concept is accepted, one could consider cyber attacks as, say, a pollutant, or, a cause of illness like a virus. I am going to choose the concept of the digital virus not because of its contemporary relevance amidst a pandemic, but because it's what technologists have called malicious code for decades.

In the current context the concept of a virus is actually in some ways unhelpful, because it could be taken as implying that the next global event threatening the lives of millions will be a cyber attack. I do not believe that to be at all likely. As I set out in a speech at RUSI in September, the world's experience of cyber attacks so far has been one of chronic, debilitating damage rather than catastrophic risk to life and limb. That is likely to remain the case; again that is a discussion for another time but the reasons for it are set out brilliantly in the 2013 book *Cyber War Will Not Take Place* by Professor Thomas Rid, formerly of this parish.

Where the virus analogy does come in handy is in thinking about the way so-called cyber weapons work. As per the title of this discussion, cyber weapons are called viruses for a reason. They look for a host. They spread. They enter the environment, and it's not always easy to predict, and sometimes even harder to control, how they will behave.

The power of the concept of cyber weapons as viruses is something we under-estimate at our peril. I know this from direct and difficult operational experience. When the NCSC was set up, as you might expect we war-gamed all sorts of different scenarios about the sort of major incidents we might face. As it turned out, the two most serious incidents we faced during my tenure involved something we didn't predict: an attack going viral by accident.

Both cases – within 50 days of each other in the summer of 2017 – involved the attacker losing control of the virus they had released and causing untold and unintended damage. The first was Wannacry. Wannacry was a virus released by the North Koreans to extort money from corporate victims in Asia. It ended up disrupting hospital bookings in Great Britain and platform announcements at German railway stations. A month later the Russians attacked a professional services firm in Ukraine. That virus, NotPetya, ended up seriously damaging a Danish shipping company, a UK advertising agency and a chocolate factory in Tasmania. What better illustration could there be of viruses in the digital environment?

***

Let us now consider, in the context of these concepts, the sort of cyber interventions open to the sorts of societies we are. I want, tonight, to propose a taxonomy to discuss this. I don't claim this to be definitive, or anything more than the start of a much-needed debate. It will be easy to criticise and pick holes in.

But I propose it anyway because we need to find a way of breaking down the discussion on 'cyber capabilities', 'cyber weapons' and 'cyber power' into manageable chunks. I may have missed it, but I have not seen the different types of interventions categorised in a helpful way, so what I offer tonight is just a suggestion for a framework for such a discussion, in the absence of one.

And we need a more sophisticated, technically informed discussion. Policymakers, both political and official, who would think nothing of offering a lay person's view of the case for an aircraft carrier, can be oddly deferential and therefore unquestioning when confronted with discussions about digital capabilities. The result can be a very simplistic conversation. At its worst, it is not far off a modern day equivalent of Sir Humphrey's famous conversation with Prime Minister Hacker about the nuclear deterrent:

"if you walked into a nuclear missile showroom you would buy Trident – it's lovely, it's elegant, it's beautiful. It is quite simply the best…it's the nuclear missile Harrods would sell you."

When I started out, I raised the UK's cyber objectives with one senior figure in Government and I got the reply "Where's the red button?". I was struck that I wasn't asked if there was a red button, and if there was, what purpose it served.

We've moved on a bit from that, but there is still, too often, very loose language reflecting a profound lack of understanding of the capabilities we are talking about. We hear assertions like "we need the best cyber capabilities in the world". Perhaps, even though most experts now rightly reject the analogy between cyber capabilities and nuclear weapons. But if cyber is an environment, are cyber weapons more akin to biological and chemical weapons, at least to digital infrastructure if not to humans? It's worth thinking about. Would we ever hear anyone argue "we should have the best biochemical capabilities in the world?".

This analogy is too simplistic in the other direction. I use it only to make a point. But it shows, I think, the need to ask ourselves challenging questions about what sort of capabilities we want, and for what purpose.

Similarly, remarks like "we should hit back hard in cyber space against those who attack us" imply that the cyber domain is enclosed, like a boxing ring, with only like-for-like capabilities deployable against adversaries. Whereas it's obvious, as I'll come to later, that there's no need to fight cyber with cyber.

Finally, discussions of cyber capabilities often get confused with a far bigger discussion about the role of technology in modern warfare: drones, automated weapons, killer robots and so on. These are very different, and much bigger, issues that are not the concern of tonight.

What tonight is about is, bluntly, the details of when it makes sense for the state to get into the business computer network attack: of hacking not just to spy but possibly to alter, disrupt and destroy computer systems.

***

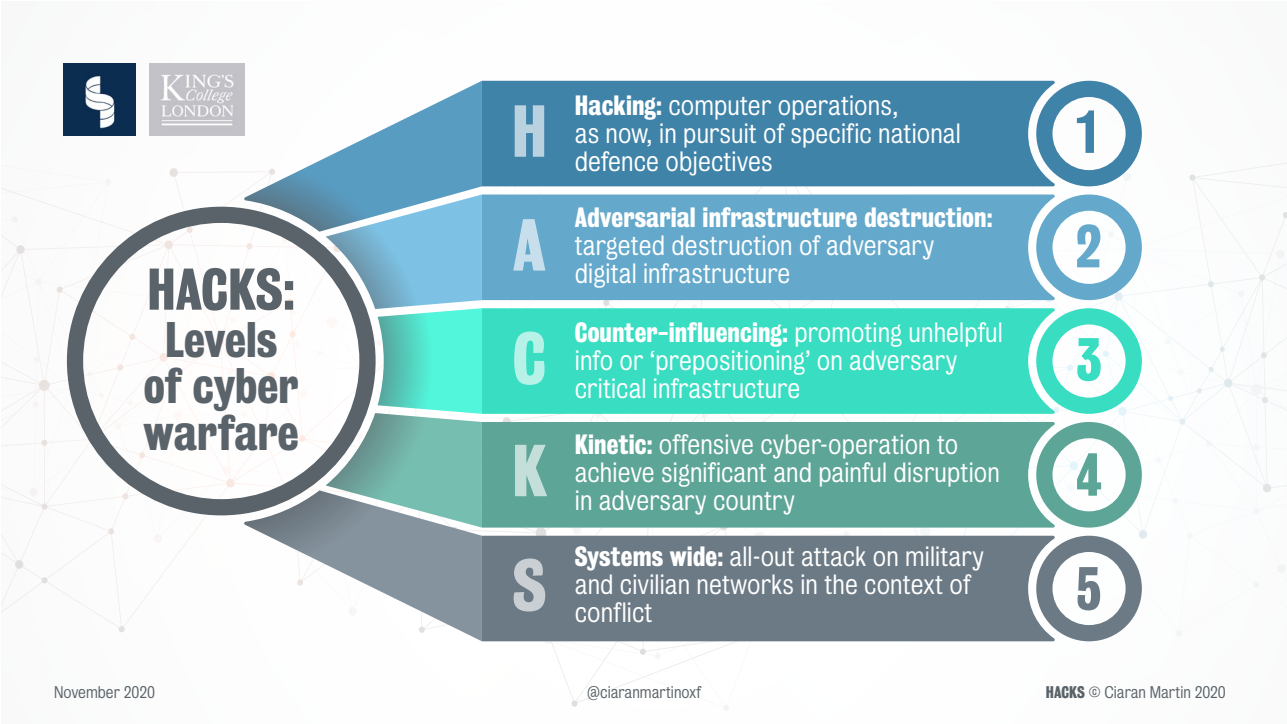

So here, for the sake of the discussion, is a five tier structure, in ascending order of severity, for thinking about the sort of capabilities nation states might deploy in the cyber domain; things we are now loosely brigading under the term 'offensive cyber':

- **Level 1:** The first level is simple **hacking** operations in pursuit of specific national **defence** objectives. In other words, there will be an intrusion on the digital infrastructure of an adversary who is acting against the state's interests by pursuing an operation against it. The intervention will degrade the actor's ability to act against the country carrying out the operation. It could be, and this is a publicly confirmed operation, destroying the so-called Islamic State's propaganda capability ahead of the Mosul offensive. It could be taking down disinformation, as apparently reported in yesterday's Times in respect of lies about vaccines. To all but the digital pacifist it is unobjectionable in principle (the targeting may be controversial, and there are healthy debates about oversight, but If it were not for the broader discussion around cyber capabilities we would not be discussing these capabilities now).

- **Level 2:** The second is what might be called **adversarial infrastructure destruction:** in other words a destructive cyber attack against *only* the digital infrastructure of a hostile cyber attack group. It is a part – though not the totality – of what the United States Cyber Command under the admirable leadership of General Paul M. Nakasone has called 'persistent engagement' on in the UK, through the NCSC and wider GCHQ, is called 'counter-cyber'. What is distinctive about this aspect of offensive cyber is its emphasis on precision in both targeting and execution. It is not concerned with sending a general 'message' to the adversary. It is concerned with destroying its ability to act. Operational details are seldom given. But it is widely reported that US Cyber Command conducted an operation ahead of the 2018 mid-terms against infrastructure used by the Internet Research Agency troll factory in St. Petersburg. This would be an example of this type of operation;

- **Level 3** can be called **counter-influencing.** This is, essentially, doing unto them what they tend to do unto us; it could be called 'digital harassment'. It would work in two ways: promoting information unhelpful to their objectives, or, secondly, 'prepositioning' on their critical infrastructure as the basis of a future attack. These intrusions would normally be covert but, as has happened in respect of Russian, Chinese, Iranian and North Korean attacks on the West, enough would be uncovered that the adversary would know that the hostile capability was there and could potentially be activated, even if they chose not to make it public. What is distinctive about this type of offensive cyber operation is that it is passive, at least initially. It does not disrupt, destroy or degrade any infrastructure;

- **Level 4** could be called **kinetic**. This is the launch of an offensive cyber-operation to achieve a significant and painful disruption in the targeted country, beyond the direct and reciprocal disruption of a specific threat. That's what makes it different from the first two levels of activity. Russia disconnecting the power supply to parts of Kiev in 2014 for six hours, turning the lights out on a quarter of a million Ukrainians, is an example of this. So too, if true, is the apparent Russian attack on the French TV channel TV5 Monde in 2015, taking it off-air for 11 hours and almost destroying it totally. There are relatively few known examples of this type of disruptive kinetic attack. But speculation about it, and its future development, are rife. There is constant talk about whether Western powers might disrupt adversary media in a time of crisis, for example. This type of activity can threaten lives at one remove, such as the removal of power supplies. But, as in the TV5 Monde attack, or the apparent Iranian attack on Saudi Aramco in 2012, it can seek to achieve highly damaging but non-violent impact too; and

- Finally, **Level 5** could be called a **Systems-wide all out attack on military and civilian networks in the context of conflict**. This does not require much explanation. But it bears a moment of pause, and of reflecting back to the concept of cyber as an environment. Because when we hear phrases, as we often do, such as "cyber will be an integral part of future conflict", this is part of what that means. We mean we would be seeking to degrade, as much as we can, the digital life of a country or area with whom we are in a state of hostilities.

So this framework: hacking; adversarial infrastructure destruction; counter-influencing; kinetic; and systems-wide all out-attack, or the HACKS framework, to give it the now obligatory acronym, is one way of thinking about how liberal democracies should think about the development and possible use of cyber capabilities.

There are imperfections and overlaps in the HACKS framework. There are also huge issues of international law – in the UK this framework starts at Level 1 with a recently settled legislative framework and ends at Level 5 under an assumed state of the Law of Armed Conflict; inevitably it's the bits in the middle that are the trickiest. I will leave those issues, for now, in the hands of the community of legal scholars working on them.

There is also, I think, two different categories of activity here. The first two categories are essentially defensive. In the first category, an individual or group is trying to do something contrary to the country's interests so you interfere with their equipment to try to stop them, or to mitigate risks, or acquire intelligence, or, in level 2, take out their computer equipment. But in levels 3 to 5 you are using digital tools for a more general projection of power: in level 3, to harass, in level 4, to hurt, and in level 5, to fight. The first two levels are symmetric countering of threat. The upper levels are asymmetric projections of strength.

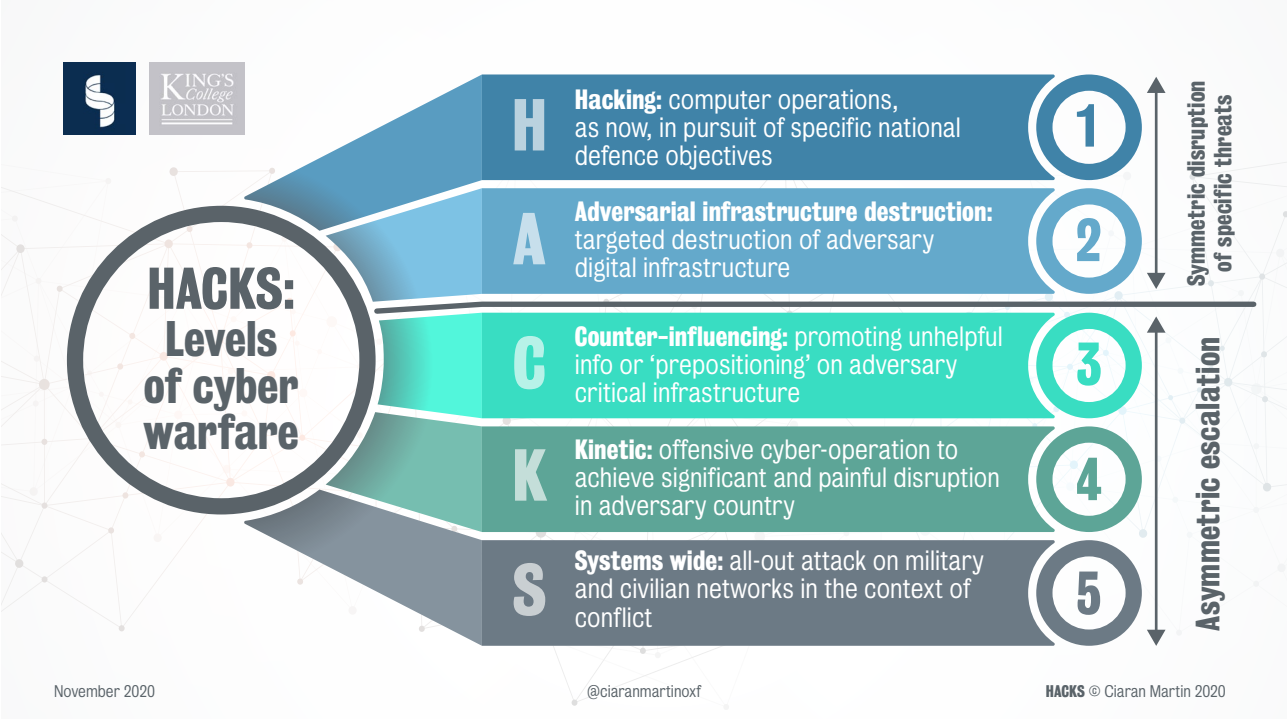

***

What I want to do now is to consider the three reasons I think Western states might want to engage in each of the various types of activity, and what recent history points to in terms of the efficacy of each of them.

The first of the three reasons why states like ours might want some of these capabilities is to defend against specific national security threats. Here, in my view, the case is most certainly made for the possession of, and willingness to use, the first level of capabilities. Put simply, we need to be able to do things like alter or delete data on a terrorist website or interfere with a hostile actor's phone in order to stop something bad happening. We need to be able to do masquerades and other operations in the digital domain. These types of capabilities are used more and more against serious organised criminals. Good.

The second reason for possessing, and being willing to use, cyber capabilities which is often discussed, is about deterring cyber attacks . Here the evidence is decidedly mixed. By that I don't mean it's inconclusive: I mean in some areas it clearly works and in others it clearly doesn't.

It seems to work in respect of the second level of activity, the direct degradation of adversary infrastructure. This is simple deterrence not so much by denial but by destruction: we have located the attack infrastructure and destroyed it so it cannot be used. US Cyber Command's operations, and those of allies, in this area seem to be achieving significant effect. If it is true, as some American media reports have claimed, that Western cyber operatives are looking at ways to take out the infrastructure used by the organised criminals responsible for the ransomware attacks that have included hospitals and other healthcare providers, I will be the first to applaud.

Where cyber attacks don't work, in my strong view, is as a *psychological* deterrent to attackers. I don't say this as a matter of philosophical conviction; I would love it to be true that cyber-retaliation deters attackers.

But it's not true.

In all my operational experience, I saw absolutely nothing to suggest that the existence of Western cyber capabilities, or our willingness to use them, deters attackers. Nor have I seen any convincing research. For more, I would refer to the excellent work of Jason Healey at Columbia University; in the interests of time I will confine myself to just one of his quotes: "After Stuxnet and the Snowden revelations, what adversaries can possibly doubt the power of U.S. cyber capabilities?. And yet, years later, the White House still complains that 'adversaries have increased the frequency and sophistication of their malicious cyber capabilities'". Quite.

And I think this reflects, strangely, this curiously confined view of cyber as a domain of operations. It implies that cyber is a confined domain where cyber must be met with cyber. We would do better to recognise that our societies are organised differently from those of our adversaries. They might irritate and weaken us with their cyber attacks because we have a free and open Internet. Our adversaries don't, so if we respond in kind we may not irritate or harm them in the same way. Instead, we can use other tools, and I am glad we are increasingly doing so. So, for example, by and large, rich and powerful allies of the British Government tend not to have huge assets in Moscow, but allies of the Kremlin do have that in London, so it is right that we move against that to achieve an effect. The new EU sanctions regime are also a terrific move by the 27 in this direction. The Obama administration's ingenious innovation of issuing criminal indictments against hostile state actors did more to deter hostile state activity than any retaliatory cyber attack: not just by embarrassing the states they accused, but by removing, for life, the prospect of traveling to the West for any of those indicted.

Finally, we must look at the development and use of cyber capabilities as a projection of wider state power. This is not about countering direct threats like terrorism, or cyber-specific threats. This is about being able to use digital capabilities, both in conflict, and in the more hybrid environment it is so common to discuss now. Therefore it relates to the more contentious, upper end of the HACKS scale.

Given that cyber attacks don't deter cyber attacks, the upper end of the HACKS scale is of no use for deterrence purposes. So what else can these types of operations achieve?

Here, the case is at best unproven, perhaps thankfully so: we have not been drawn into anything close to cyber conflict. But there are good reasons to be cautious about the efficacy of these capabilities and realistic about what we can hope to benefit from having them.

Levels 1 and 2 work to counter specific threats. So let's consider the third level: counter-influencing through digital harassment. One part of this toolkit – information operations – is as old as statecraft. But we will want to be careful not to be seen to be as willing to undermine others' political discourse as they are ours. More interesting is the case for 'prepositioning' – intruding upon a hostile state's critical infrastructure, as they do ours, for leverage, and potentially disruptive action, in the future.

In 2019, the New York Times reported that the US had sought to penetrate the Russian electricity grid. It was met with an angry denial from President Trump, accompanied by briefing from his Administration questioning why the US would be interested in attacking such a network. I have no idea about the underlying truth, and couldn't comment if I did. What it illustrates, though, is that this type of activity is only a credible operation if the state is willing to contemplate following it up with an actual attack.

And are we? Is an attack on the power supply of another country something we will be willing to contemplate? Will we actually even take out a broadcaster, even if we could? In what circumstances would it be lawful? Does the other country see it as a serious proposition? We need to think those questions through fully. And for now, I suspect there is sufficient doubt over this to cast a question over the credibility as a projection of state power.

In a sense, moving to the highest level of capability for use in all-out conflict is easier to think about in this context. But the real question here is what sort of capabilities do we actually mean, and whether we actually understand how they would work. Yet again, this comes back to the virus analogy: in these circumstances we would be releasing very destructive, very potent capabilities into the digital environment. Are we sure we can target them properly? Are we confident they can't just be caught and fired back at our own civilian infrastructure? If they are, can we protect against them?

So there are some pretty significant questions about the *efficacy* of the higher range of offensive cyber capabilities. But we should also be particularly cautious about the efficacy of the upper end of cyber capabilities because of the potential risks such cyber weapons – viruses – carry for our own digital environment on which we so depend.

***

There are two tactical and one big strategic risk for open, democratic states in developing and relying heavily on the higher end of offensive cyber capabilities for strategic national security.

The first tactical risk is securing the capabilities. It is an uncomfortable fact that the record of Western security services in guarding cyber capabilities is not perfect. There have been thefts and leakages, and this has put some knowledge of malicious cyber code into the public domain where less responsible actors can and have used it against our interests with consequences for ordinary citizens.

These capability breaches are not, in my view, solvable simply by ordering those in charge to do better at guarding them. It is in some respects an inevitable risk when we are talking about capabilities that are designed to be used on linked-up computer networks. It is very hard to develop them and store them with 100 per cent security 100 per cent of the time.

In my view, it is irresponsible for Governments to plan on the basis that they can develop and store cyber capabilities on the assumption that they will never leak or be stolen. So this aspect of risk management must be priced in when developing these tools. No one is likely going to be able to steal a nuclear weapon. No one will accidentally lose or leak a ballistic missile. A fighter plan cannot be stolen without someone noticing. None of these statements hold true for cyber capabilities. And the risk becomes more potent the higher up the HACKS scale we go. By definition, these are capabilities that we will want to use sparingly, if at all. So how do we ensure their security at rest?

Secondly, there is control of the capability once used. I mentioned already that the two worst incidents in my career were state attacks gone wrong. Put simply, what if the West did a NotPetya? What would we think if we turned on the TV and on the news was chaos across corporate Asia (for sake of argument) because a Western operation had gone viral? Are we sufficiently confident in our own operational infallibility to believe it will never happen? If so, why? Furthermore, once the weapon is out there it can be studied, reverse engineered, and used again. What do we do with that?

Then there is the strategic challenge around balancing offensive and defensive capabilities. And here we come back to the concepts of what it means to have digital viruses in our digital environment.

Our societies will never be the winners from insecure technology and an unsafe Internet. Authoritarian, less digitised countries will be more reckless in exploiting weaknesses and more tolerant of pain in the event of escalation.

Therefore, we must be unambiguously in favour of safer technology. That holds even if that sometimes makes deploying our own offensive cyber capabilities harder because a rising tide of security will, to some extent, lift all boats, including adversarial ones.

And this is not really a moral argument: it is one of efficacy. The practical risk calculation for a developed, open market, rule of law democracy is clearly in favour of prioritising defence. For a much fuller and technical exposition of that argument, I would commend to you a July paper in the Journal of Cyber Policy by my former Israeli opposite number and good friend Dr Eviatar Matania, and his colleague Eldad Tal-Shir.

For me, what tonight's analysis shows is that the efficacy of cyber capabilities is proven at the lower end of the capability scale. Cyber effects can and are being used to counter terrorists, serious criminals including online child sex abusers, and propogandists, hackers and malignant online actors of all kinds. But the threats they are countering are about public order, public safety, or digital public health, if you will.

They represent many unpleasant things. But this is not war.

Thinking of it as war is playing the adversary's game. Just because authoritarian states have extended their harassment of us into the digital domain doesn't mean we should, or have to, automatically accept the weaponisation of cyberspace. These days, the question is often asked: "what does the cyber domain mean for warfare?". There is not enough attention paid to the question the other way round: how do we minimise the impact of digital harassment by adversaries on the domain of peaceful social and economic activity that is cyberspace?

It is, to me, far from proven that escalating tensions in the so-called grey zone is an effective way of doing this.

***

I recognise that many will take a different view. That is at it should be and there should be an open debate.

But I worry that there isn't one. We risk an acceptance that the acquisition and use of higher end cyber capabilities are a priority, without testing the question about what this means for our own digital environment. We haven't had this fundamental debate because the national security community and the technological communities are not really talking to each other. We risk being modern day Sir Humphreys in the Harrods weaponry showroom.

I do welcome the steps to shed some light on how we think about these emerging capabilities. Both the UK and the US have disclosed some details of operations. The White House and GCHQ deserve great credit for setting out, respectively for the US and UK, the process for deciding whether or not to make a flaw in technology public or keep it secret to use in security operations. But it would be easy, and confidence building, to go further. Why shouldn't we disclose more operations, given that for now at least they are rightly focussed on crime, terrorism, propagandists and cyber criminals? Why can't Governments say more about how they manage the risks of their cyber capabilities being stolen, or of attacks going wrong? What are the oversight arrangements to ensure good risk management of cyber capabiltiies?

Surely being more open about the sorts of things we do under law before a crisis is the great lesson of the Snowden affair?

And isn't the lesson of the successes of cyber security in the UK over the past few years that transparency is not just more possible than we thought it was, it is hugely advantageous because sharing information about risk helped people manage it?

***

Trust is everything, and trust can't be built in secret.

But I am just calling for a more transparent debate. I have a very passionate view about it.

So let me conclude by restating my strong view that we must start from a robust assertion of three principles about our place as open, liberal democratic societies in the cyber domain and the cyber environment.

The first is the primacy of cyber defence.

The second is the pursuit of a safer Internet and safer technology as a clear policy goal.

The third, flowing from the first two, is that in developing offensive capabilities, and contemplating their use, liberal, democratic nation states should exercise great caution, particularly the further up the scale they go. In the foreseeable future, their efficacy is in promoting digital public safety and security in support of a healthier digital environment.

***

If we get this right, there is a real opportunity to make us more truly secure in cyberspace.

We know what we have to continue to do. We need to build genuinely resilient systems that won't fall over because of one attack. Fix hardware. Promote safe software. Help stimulate the cyber security market so it's easier to buy perimeter defences and internal monitoring systems that can fend off attackers. Train people better, and upskill the next generation.

And, crucially, we need to correct the mistake of the last generation of technology by getting in ahead of time as newer technologies like the Internet of Things, 5G, AI and quantum to build in security and resilience from the start.

We must focus on promoting security in technology by design. And we have to convince people we are serious about promoting a more secure digital environment; the absence of trust that we are is one of the reasons we struggle to get a hearing on issues like encryption, for example.

And if we get this balance of capabilities wrong, and there is a real risk that we might, I believe we will end up less safe.

We will have weaker, less well defended citizens. They will be ever more dependent on vulnerable and unsecured technology. And we will be trying to counter this by relying on deterrence capabilities that don't deter, grey zone capabilities that might backfire, and top-of-the-range capabilities that we will never use.

Perhaps most tragically of all, we will have allowed more authoritarian regimes to goad us – to rise to their bait – into securitising what is a great creation of our Western culture of freedom.

Of course we should reserve the freedom to act in the cyber domain when we need to. One of the first things I did in office was to sign, in public, a Court statement defending the lawful use of cyber capabilities to combat threats to the UK. It was a lonely place to be, in the aftermath of the Snowden crisis. But it shows I am not a digital pacifist: I believe strongly in the right to take action when necessary.

So in arguing for restraint, I am not arguing for weakness.

The case for cyber restraint is a hard-headed one.

A more secure digital environment is the best guarantor of safety and security for Western countries in the digital age.

We weaponise the Internet at our peril. We militarise the Internet at our peril. We should show we have the courage of restraint to protect its magic.

In the cyber domain, the best form of defence is defence.

Ciaran Martin is Professor of Practice at the Blavatnik School of Government at Oxford University.

Until August 2020 he headed up the UK Government's National Cyber Security Centre (NCSC), which he established as its first CEO in 2016.

Prior to that, Ciaran was Constitution Director at the Cabinet Office from 2011, working on the Scottish independence referendum.

From 2008-11 he was Director of Security and Intelligence at the Cabinet Office.

@ciaranmartinoxf

