Cyber Attacks: What actual harm do they do?

Ciaran Martin, RUSI, 18 September 2020

There is an essay by the American satirist PJ O'Rourke called *Studying for our Drug Test* (it is reprinted in his 1992 Collection *Give War A Chance*; probably some way down the list of choices for RUSI's next series of seminars). In it, O'Rourke describes an announcement by the Reagan administration during its Wars on Drugs, issued this very week in 1986. The proclamation read that every federal employee would be subject to randomised drugs tests, along with millions more private sector workers in sensitive jobs involving public safety and security.

Much of the accompanying fanfare focused on the strict new testing requirements for anyone working on board a commercial aircraft. So, as a sceptic of Government intervention, to put it mildly, O'Rourke rang the Federal Aviation Administration. He asked: "How many fatal accidents on major airlines have involved drug use by flight crew, air-traffic controllers and other responsible personnel?"

"The answer is none" replied the FAA spokesman.

This anecdote came back to me as I reflected on my six and a half years running cyber security for the UK Government, the last four of which as the proud, founding chief executive of the National Cyber Security Centre. (I can't continue in my first public speech since leaving the NCSC without paying tribute to the brilliant team there and the partners who worked with us. We did some good, and I hope at a time when people are, quite rightly, asking hard questions about the state's ability to grapple with strategic threats, the NCSC showed that expert, imaginative, innovative and dedicated public servants can make a real difference in the right environment. More on that some other time. For now, back to PJ O'Rourke and his drug-addled pilots.)

The reason I thought of this story was when I looked back to my first public appearance as GCHQ's awkwardly-titled 'Director General for Government and Industry Cyber Security' (the rationalisation of the job title alone was reason enough to set up the NCSC). I was due on stage at a skills competition for young people. I was to be preceded by the obligatory video. I was nervous: new to the subject and back in those days, after years in the shadier parts of the civil service, genuinely unaccustomed to public speaking.

My experts told me that it would be good to mention that it was important not to overhype the cyber threat, so that we could get a more realistic appraisal of the problem and apply appropriate technical and organisational expertise to it. That was the bulk of what I thought was my safe, uncontentious, technocratic and frankly boring speech.

Imagine my surprise therefore, when the stage fell dark and some scary music started. Then the video started. I saw burning streets. Through the flames emerged angry rioters. Hungry looking people were queueing at empty ATMs; others were looting shops. And so on. You get the picture.

Once I had mangled my way through a singularly inappropriate speech to follow that video I began to think of this as my PJ O'Rourke moment. Let me explain.

One of the points of the O'Rourke essay is that you don't help find good public policy solutions to complicated problems by way of hype with no basis in evidence. It was inconvertibly true in 1980s America that there was a serious drugs problem causing all sorts of chronic, devastating social and economic harm. It was also true that catastrophic aviation failures were more common than, thankfully, they are now, and that anyone taking a flight had to contend with the small but observable risk that they might be put in a situation on a plane where something would go wrong that gave them and large numbers of other people no chance of survival.

The point was that there was literally absolutely no connection in evidence, none whatsoever, between the two problems. By linking one serious problem – drugs – with a major catastrophic risk – plane safety – the US Government wasn't helping to tackle either of those problems. And by making such a high-profile announcement, it was probably hindering the public understanding of both problems, and increasing public fear without any basis for doing so.

Clearly the policy in and of itself was not objectionable. Even if drugged pilots were not an actual problem, no one would say they are a good idea. Indeed, whilst we will never know the counterfactual, given that most countries don't allow pilots to take drugs, it might have prevented some harm from happening in the future even if it had never happened in the past.

But there are two clear problems for the public policy maker here. One is displacement. Policymakers have limited levers: limited time, limited powers, limited ability to get the law changed, a limited public attention span if you want to change people's behaviour. If the aim is to either tackle drug use, or enhance aviation safety, this measure would not have been at the top of anyone's list. And if another aim is, as it should be, to enhance citizens' ability to take good, risk-based decisions about both the drugs menace and flying, this measure did not help either.

Let us turn this lesson upon cyber security. Right up front, let me say that a cyber security catastrophe is a possibility, and has been a possibility for some time. I am not saying it is something that should be ignored and will say more about the long-term outlook a bit later.

What I am saying is that after nearly three decades of warnings, a catastrophic cyber event has yet to occur, in terms of the normal definitions of catastrophic events.

Yesterday evening, unconfirmed reports emerged from Dusseldorf in Germany of a woman, who suffered from a life-threatening illness, dying *en route* to hospital. Police are investigating, but some press reports say the cause of death was an overly long journey because the nearby hospital which was supposed to receive her could not treat her because its systems had been crippled by a ransomware attack (a ransomware attack is something I will explain in more detail later in this talk; for now, suffice to say it is usually carried out by organised criminals).

If true, this tragedy would be the first case I know of, anywhere in the world, where the death of a human life could be linked in any way to a cyber attack. Until it is confirmed, there remain no proven cases of fatalities from cyber attack. And if it is confirmed, it is highly salient that it is the inadvertent result of amoral and utterly reckless cyber crime, not premeditated state or terrorist attack designed to kill or hurt someone.

So the predicted catastrophes and cyber wars have not come to pass. But all sorts of other harms have been caused by cyber attack. Skewing our focus onto the catastrophic risk in spite of the evidence, which certainly was the case when I started six and a half years ago, risks skewing the right policy response.

This isn't pointed out often enough. I had rather hoped to use this lecture be the first figure in the cyber security debate to point it out for some time. Sadly for me, I am now John Landy to the Roger Bannister figure of the excellent Jim Lewis at the Center for Strategic and International Studies in Washington, who published an important article on this exact subject a month ago. To quote his opening paragraph:

"As a trope, a cyber catastrophe captures our imagination, but as analysis, it remains entirely imaginary and is a dubious basis for policymaking. There has never been a catastrophic cyberattack".

I'd argue that it's not just a dubious basis for policy making; it is the wrong basis. The first task of anyone trying to set a course for good public policy and corporate and individual behaviour on cyber security is to make a realistic assessment of the problem. As you might expect from someone who has just joined the Blavatnik School of Government, I believe in evidence-based policymaking.

But today I am not going to make policy recommendations about cyber security. I am simply going to tell you what I saw in my six and a half years running a national cyber security effort on behalf of the UK. So it's not scientific. It's, inevitably, a very Western view of cyber harms and we should seek other voices from across the world to listen to cyber harms from their perspective. And my categories can be debated, and they're imperfect and probably overlap a bit.

But this is simply what I saw, and a few lessons from it. Here are my observations about the types of cyber harms.

Getting robbed

I start with 'getting robbed' as the first of three categories of manifest cyber harm. This is because this category includes some of the most common cyber attacks, and it is the simplest to understand.

First, there is have straightforward cash theft. The stealing of money. This could be anything from the small amounts skimmed through cyber-enabled fraud to the eye-catching antics of more sophisticated hackers, who, in February 2016 targeted the endpoints of the Swift interbank clearing system and transferred US$81m to accounts in the Philippines. It appears that a small spelling mistake prevented the extraction of up to $850m more. A number of countries and cyber security companies have blamed this attack on North Korean state-backed groups and it is certainly the case that this cash-starved dictatorship has become the first nation state trying to steal actual cash through cyber attacks. At the other end of the spectrum, 'routine' cyber crime is notoriously hard to measure but, to give one of many examples of contributory data, the Action Fraud hotline reported cases it knew about in 2018 which amounted to the recorded and verifiable loss by British citizens of £190,000 per day, and this is presumably only a fraction of the real total.

The second part of getting robbed is more complicated and more commercial: the theft of intellectual property. In my experience, corporate-on-corporate hacking of IP is much less prevalent than most people believe (certainly in Western countries where the legal consequences of getting caught spying on competitors are severe). This is more of a nation state problem. One of the many ways in which Chinese and Russian state sponsored cyber threats differ (despite, often unhelpfully, being lumped together in cyber security parlance as if they were the same threat) is that the origins of Russia's cyber objectives are political whereas China's are economic. Way back in the first decade of this century, ubiquitous and often noisy Chinese cyber attackers hacked corporate networks and stole designs of high value products; suspiciously cheap and quickly-designed home-grown competitors would appear on the Chinese market thereafter. Since the (separate) 2015 agreements between the US and the UK and China on commercial espionage, this activity has been quieter but the December 2018 attribution by the UK, US and Australia of a huge, global campaign of hacking by Chinese state sponsored attackers (known in the trade as 'Cloud Hopper), which compromised many of the main IT services providers and through them their clients, showed it was a threat that hadn't gone away. Cloud Hopper – described by FBI Director Christopher Wray as affecting the 'A to Z of American business' – also shows the difficulty of quantifying the damage. The way the attack worked (and some poor security on the part of victims) gave the attackers access to at least some of the clients of the primary victims; thanks to poor monitoring by some of the primary and secondary victims we can never fully know what the attackers got.

Finally, we are getting robbed of data. Lots and lots of it. The lists are long and well known. 46 million Target shoppers in the US. Equifax. British Airways and Marriott Hotels, who have been subject to huge regulatory fines in the UK.

In the spirit of trying to understand the harms caused by cyber security, however, it is worth pausing to reflect on two related complexities in the problem. One is about what actual harm is being done; and the related challenge is the severity of the harm.

Personally, I don't like the way data breaches are presented in the media in terms of headline records lost. That's because it is only one of the indicators of harm. Let us consider two hypothetical cyber attacks. In the first, it is reported that 100 million personal records have been stolen. But those records are just names, email addresses, and the last four digits of a bank card. In the second, it is reported that half a million personal records have been stolen. But they are, for sake of argument, the entirety of the content of mortgage applications or any form with detailed identification requirements, with not just primary information like National Insurance numbers, but common secondary identifiers like mothers' maiden names. In the latter case, those affected are at far greater risk of suffering direct harm through identity fraud than in the first case. Indeed, I would argue that more harm is done in the second attack, despite the headline number of 'victims' being smaller by a factor of 200. It is nuance like this that understanding of cyber harms currently lack.

Another area where understanding needs to be deepened so that better interventions can be made is to understand why hackers target datasets. The EU's General Data Protection Regulation, which as anyone who ever tries to read an American newspaper online knows, has significant effect beyond EU countries, permits heavy punishment for those who fail to protect data. But what harm is it punishing? Most of us can find ourselves on the brilliant website **haveibeenpwnd.com**, to see if we were hit in any one of the big data breaches of the past. But if we are, other than a deep sense of unease that someone has our personal information, what harm are we at risk of?

The best account I've seen of this remains the NCSC's 2017 paper mapping out how an organised crime syndicate works. That paper sets out, amongst other things, how stolen credentials can be sold on to other gangs, matched against other data sets to provide for better and more-targeted identity fraud or other criminal exploitation, or provide the basis for a brute force attack on a different organisation because of the common practices of reusing important passwords and failure to apply multi-factor authentication. Understanding the workings of global cyber attackers in this way helps: it means that we don't panic people into thinking that their bank accounts are going to be emptied when they get notified that their personal data has been stolen, which for the most part isn't normally the case (for example, there was, once the full facts emerged, no need for the sort of panic we saw in October 2015 when TalkTalk lost tens of thousands of largely outdated, minimal records). Instead, we should focus on public policy interventions that break up and disrupt this criminal cyber attacking ecosystem.

One large scale data breach where both the intent of the attacker, and the harm caused, are relatively easy to understand is the Chinese state attack on the Office of Personnel Management, confirmed in 2015. This attack is believed to have resulted in the theft of the personnel records of 22.1 million past, present and prospective federal Government employees in the US. But this wasn't passwords and incomplete credit card numbers – this was the detailed security clearance forms with all family members, college roommates, behavioural preferences, and every immediately obvious personal record, and also, it is believed, 5.6 million fingerprints.

Getting weakened

The OPM case is an illustration of serious and direct harm through cyber attack: US national security leaders at the time were open about that. It meant that not only did the Chinese state have what the then Director of the FBI called "a treasure trove of information about everybody who has worked for, tried to work for, or works for the United States government", but it also meant that the people affected realistically had no trustworthy method of secondary authentication available to them, because Beijing had either some of their fingerprints, or their mother's maiden name, or whatever.

So I will actually put the OPM attack as an example not of getting robbed, but of the second main category of attacks I have seen in my operational service: getting weakened.

By this I mean cyber attacks which, although they may not cause direct physical or even economic harm, strategically weaken a nation and a society. As such the perpetrators tend to be (as in the OPM case but this isn't always the case) nation states, and the incidents tend to attract the interest of national Governments. The UK Government has, since 2015, talked the language of 'strategic advantage' when discussing national security strategy. These harms are examples of when the state is strategically disadvantaged by dint of cyber attack.

I have seen three aspects of national interests getting weakened through cyber attacks. The first is obvious and as old as the hills: espionage. Electronic spying is now a well-established practice. Foreign Ministry networks are one of the most targeted of any Government. Any major international political event risks becoming an international festival of electronic espionage. And the recent attempts – thankfully, so far as I know, unsuccessful – to breach the UK's vaccine research is an example of the ongoing problem.

Many scholars of espionage will see this as business as usual via electronic means, and there is some truth to that. Why I keep coming back to the OPM case is that it shows the potential for systemic and strategic harm, rather than just the sort of short-term tactical advantage espionage often confers. Electronic espionage can be done at scale: if the average traditional paper vetting file weighs about half a kilo, and I can confirm without breaking the Official Secrets Act that it probably does though I won't get into the specific of either mine, or my host's, then for a manual spy to infiltrate the OPM to the same effect would have required 11,000 metric tonnes of paper, amongst other things.

The second weakening arises from attacks that weaken reputation and confidence in societal discourse. There is here a bit of a private sector angle with so called 'hacktavists' seeking to mess around with corporate reputations by doing things like hijacking Twitter accounts or defacing websites to embarrass the company. But the primary strategic aspect of this is interference in a country's political discourse. In my view, we talk too much about election security and not enough about protecting the political system as a whole: lived experience shows that this threat doesn't just pop up at national election time. To prove the point, the first case in which the UK has confirmed an attempt at outside interference from Russia in the political process relates to the Department of International Trade papers relating to the

National Health Service. Publicly available evidence shows that this activity continued right throughout 2019. All of us who remember the politics of 2019 will recall that there was no certainty about whether there would even be an election at all that year, and if so, when it would be.

One further comment on political interference. I think it is clear that strategic harm has been caused by the Russian campaign against Western political discourse: the evidence is the very fact we're talking about it. It shows the attackers have managed to inject a scintilla of doubt into our societies and that is concerning. But, as always in cyber security, countering this requires facts, evidence and transparency where possible, and calm reflection on those facts, rather than hype. Yes, there is a clear problem here. The UK has published its account so far about the events of 2019 and a criminal investigation is underway. Much has been said and written about what happened in the US in 2016 and in France a year later. But it does us no good to overhype the adversary, or to imply damage where none has been caused. So, for my part, let me repeat that I personally never saw any evidence of attempted interference, successful or otherwise, in the 2016 referendum. And as for the Scottish referendum two years earlier, I am somewhat baffled by the renewed speculation about it, because there appears to be no credible evidence of any serious campaign by Russia that could have influenced either the casting or counting of a single vote in either direction.

Our democratic processes are at risk of strategic harm from outside interference, but they're also much more robust than they're often given credit for, and it's in our interests to say that and retain public confidence in them.

The final example of getting weakened – of strategic disadvantage – that I will describe is so-called 'prepositioning'. This is when an attacker – normally on behalf of a nation state given what they're doing – puts down an implant on a strategically important network. It is a landing point for further exploitation. It might be there for espionage purposes. But it might also be lurking there so that, if the need arose, it could be the basis for a destructive attack against the critical network. When one discovers such an implant, it is rarely possible to be certain why the attacker has put it there, and it might be for both purposes. The 'harm' here is that the organisation and the national interested is weakened because a foreign power has leverage they could exploit should they need and want to.

Examples of this are many: indeed one of the great reforms of the last few years has been a willingness by Governments to publicise the details of such intrusions and release technical advice so that organisations can clear them off their networks. Energy and telecommunicaitons companies have been amongst the worst hit.

Getting hurt

Prepositioning is then the basis for those occasions where getting weakened turns into the final category of cyber harms I have observed as an operational leader: getting hurt. What is specific about these attacks is that they are not designed to steal or spy or influence: they are designed to stop a computer network system working, or to alter its workings, as Stuxnet apparently did, and cause harm by doing so.

Again, I will refer to three categories of this type of harm.

First is a destructive attack on something of critical importance. Although, as we have seen, no one has yet been deliberately killed by a cyber attack (and it is still possible that no one at all has been), there are several examples of serious harm being caused deliberately or accidentally. We all know that in Estonia in 2007, banks, media outlets and Government agencies were swamped with spam and malicious code and key services stopped working. Eight years later a quarter of a million citizens of Kiev were left without electrical power for between one and six hours following a Russian attack on the Ukrainian power grid. The accidental destructive attacks of Wannacry (by North Korea) and NotPetya (by Russia) within 50 days of each other in the summer of 2017 did combined economic harm of more than $12 billion. In 2012, Saudi Arabia's national champion Aramco suffered serious commercial damage via operational destruction thanks to the likely Iranian attack which deleted 32,000 hard drives. The mysterious Mirai botnet attacks of 2016, which ingeniously hijacked a range of Internet of Things physical gadgets and pointed them at an Internet backbone company called Dyn, knocking it over and taking out Twitter and Amazon on the east coast of the US for six hours, is another example and, interestingly, an example of a fairly obscure area of critical vulnerability in the Internet age that most people haven't heard of.

These were deliberate, targeted attacks (initially at least) because Wannacry and NotPetya subsequently spiralled out of control. But the second category, and one of increasing concern to me, is when the disruption risks significant social and/or economic harm, but that isn't the point of the attack. This brings me to ransomware.

I have put ransomware as a disruptive category in its own right because it is such a huge problem right now. Ransomware works by stopping a system working by encrypting the data of an organisation and demanding a payment to decrypt the data. Therefore, ransomware is theft but it's theft by extortion, and the extortion only works because the attacker has crippled the system. That's why ransomware needs to be treated as a disruptive threat, not like data theft or espionage.

Right up until my last hours at the NCSC, I remained of the view that the most likely cause of a major incident was a ransomware attack on an important service. For the attacker, the choice of the service would be incidental: they were just after money. But from the point of view of national harm, that incidental choice of victim could be important.

And again, to my last hours at the NCSC, what most kept me awake at night was physical harm inadvertently resulting from ransomware. Attacks on healthcare providers in Germany and the Czech Republic at the height of the pandemic were very scary. Sadly, it appears that worse may have happened in the last twenty four hours, though we await the full details. In any case, some researchers have begun to publish tentative evidence linking ransomware attacks on hospitals with poorer medical outcomes, including mortality rates.

Criminal ransomware, used recklessly by amoral criminals, is one the biggest but least discussed scourges of the modern Internet. It is the most likely cause of disruption of key services. It is undeniably a huge source of financial loss. It is the most likely way someone is going to suffer serious disadvantage, or get hurt, or even get killed, which may, sadly, have just happened for the first time.

This judgment is based on real, lived experience. Take two incidents over the past two years.

The first is the attack on a company you'll likely have never heard of – Eurofins, a Brussels based forensic services provider. When it was hit by a ransomware attack, much of England and Wales's criminal forensic service capability was disrupted and only brilliant work by a range of partners, mostly in law enforcement, avoided some serious public order consequences. Similar, earlier this year the borough council of Redcar and Cleveland suffered a major attack which cost is some £10.4m and left 135,000 British citizens with severely disrupted access to essential services for some time.

I pause on the Redcar and Cleveland attack because it is in countering threats like this that at least some of the attention needs to be. Criminals blackmailing English local government is about as far from the Holywood apocalypse vision of cyber attacks as it's possible to be. But these are our schools, our services to vulnerable people, our environmental protection, all at very real risk from fairly common techniques and tools. It is this problem that we need to shout about to help people understand it, and get after it.

There is one slot left on my grid to make the perfect 3*3 matrix. I will allocate it here, but with an asterix. That's because this is the one harm which has yet to happen. A hostile attack resulting in fatalities.

It is worth discussing briefly why this might be. Like Jim Lewis, I think there is a fairly simple but important answer, which, thankfully, holds true at least for now.

The bad news is that causing disruption, pain and economic harm through cyber attack, and even putting small numbers of people indirectly at risk, as we have seen with ransomware, remains too easy for my liking. The better news is that killing large numbers of people by cyber attack remains thankfully quite hard. The capabilities to do it are in the hands of only a very small number of nation states, and it is currently not in the interest of any of them to use them, any more than it is to fire live rounds at their adversaries. To put it in actual Hollywood terms, a teenager can probably hack into some system or other which belongs to the Pentagon or the UK Ministry of Defence. However, although Matthew Broderick's 1984 film War Games is a magnificent film, a sole operator from a suburban bedroom cannot launch a nuclear missile via cyber attack by guessing the password.

Let me be clear. Serious harm is possible in cyberspace. In my years at the NCSC I came close to declaring a Category 1 cyber attack – our most severe level. Category 1 didn't and doesn't automatically equate to catastrophic threat to life but it would amount to a serious national problem. If, for example, the OPM attack had happened in the UK, I would have ranked that as a Category 1. Had WannaCry started to impact critical patient care, that would have been a Category 1. And had the Chinese attack on the IT companies yielded more direct evidence of serious economic damage, it could have crossed the threshold. Moreover, some of the more severe disruptive attacks could have had wider implications had the circumstances been different: if a supplier to the Armed Forces is disrupted at a time of conflict, for example, even if that wasn't the intention of the attacker.

And catastrophic harm, at the top end of Category 1, *is* possible. Just because cyber catastrophes haven't happened yet doesn't mean they won't in the future, either via existing or new technology like AI, or that some of the measures put into place have helped prevent them.

Cyber Attacks: Categories of Harms Caused

 **Getting Robbed**	 **Getting Weakened**	 **Getting Hurt**
Cash theft	Espionage	Destructive
IP theft	Political Interference	Ransomware
Data theft	Prepositioning	Catastrophic*

But my concern about overly focussing on the catastrophic risk is fourfold.

First, it skews resources. If it's all about catastrophic risk then the tendency is to focus Governmental attention on military asset protection and the 'hard' infrastructure that is obviously critical. But in real life, the failure of something called a Domain Name Service company, or a ransomware attack on a local authority, or, worse, on a healthcare services provider, can have severe consequences. And we learned from the threat to electoral processes and from the pandemic that critical infrastructure is also local electoral administrators and hospital gown providers as well as the Army and the National Grid.

Second, it skews the attention of policymakers. As discussed earlier in the aviation and drugs context, policy levers and policy capacity are finite. A focus on catastrophic risk probably leads policymakers towards some interventions like mandatory box-ticking compliance exercises for national security bodies to show preparedness for cyber catastrophe, or working out what the threshold for various cyber retaliations should be. These could be useful, in and of themselves.

But if I had one policy card to play in the next year, I would ask for a serious examination of whether we should change the law to make it illegal for organisations in the UK to pay ransoms in the case of ransomware. The case is not a slam dunk, and if the answer is no, then we should think of something else to counter ransomware, the single biggest contemporary scourge in cyber space. At the very least, it is worth looking at a curious anomaly arising from the fact that our extortion laws are largely based on the experience of kidnappings by international terrorist groups. So if you are ransomwared by a proscribed terrorist group, it is illegal to pay, but if the attackers are what those of us who grew up in Northern Ireland in the 1980s learned to call 'ordinary decent criminals', or even state attackers, then it's fine. Surely that needs a look.

Thirdly, catastrophic risk terrifies people and that terror leads to helplessness, and helplessness is terrible for cyber security.

If the popular perception of cyber attack is food shortages, empty cash machines and riots on the streets, the average person, and the average organisational leader, will assume he or she can do nothing about it and it is problem entirely for the state, and maybe the board of a large business. This is not true. If however, the cyber risk is accurately understood, everyone knows that patching the boxes on a network and doing security updates on a phone is a very valuable thing everyone can do.

Finally, focussing on the catastrophic threat doesn't even help manage catastrophic risk that well. It deflects attention away from some of the generic things we can do in cyber security to counter the everyday threat, which also build in resilience to critical systems. Sadly, in aviation security regulators have had to learn how to contain a malevolent presence on board a plane and try to minimise the damage someone or something that has evaded pre-board checks can do. Similarly, good cyber security means assuming someone truly evil gains access to the network: how can the defender minimise the harm they can cause? How can a system fail safely? And, at a more strategic level, how do we fix the next generation of technologies to ensure they don't have the structural security flaws of the first? How do we ensure that the age of the Internet of Things, artificial intelligence and quantum computing don't have in-built security flaws? We do that by designing and regulating them well now, not by speculating on their potentially apocalyptic consequences.

Conclusion

These lessons will inform some of the themes to which I will return in future work, alongside others like: the motives of attackers; the role of economics in cyber security; the role of offensive cyber; the changing geo-political balance in tech, and much, much more. As someone new to academia let me get out of the way quickly my first statement that more research is needed.

The purpose of this talk has been to show that lived experience in the UK, and those of the UK's allies, of the cyber threat demonstrates that understanding the nuance, complexity and detail of the sorts of attacks we have lived through is crucial to making our digital homeland safe.

That lived experience shows that whilst the catastrophic risk remains, right now, cyber attacks are more a threat to wealth than to our safety, more a risk to our sense of liberty, happiness and wellbeing rather than life and limb. Quantifying the costs of cyber attacks is hard – some data losses (like those of Her Majesty's Revenue and Customs in 2007) end up causing no observable harm, whilst the Chinese attack on IT companies is almost unquantifiable. And in commercial espionage, views will differ as to what extent the commercial loss to a company equates to national harm. And so on. That said, more data on harm would be very useful: the UK last tried to compute the cost officially in 2012, and perhaps a renewed effort to get internationally-comparable data would be timely.

What is derivable is a sense of how the harms happen. The key conclusion I draw from my years leading some brilliant people with world-class capabilities and having the opportunity to study properly for our cyber security test, is this. Unlike the world of military capabilities, the digital domain is now part of our everyday environment and our everyday experience. So everyday risks are there, and those risks materialise as everyday harms every day of the year. In aggregate, they add up to a significant national security and prosperity problem. Some specific attacks weaken us at national level. Some may even hurt us, and in ways and through means we didn't predict.

So the solutions must be based on an overall levelling-up of our digital security. Governments and lots of business organisations need a lot of security. But everyone needs some security. 'Patch Squad' may be as far from the Hollywood version of cyber heroes as it's possible to imagine. But it would be a very good thing.

Professor Ciaran Martin, CB

University of Oxford

September 2020

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

