

Chapter 2

Who is the Opponent?

Going all the way back to early time-sharing systems we systems people regarded the users, and any code they wrote, as the mortal enemies of us and each other. We were like the police force in a violent slum.
– ROGER NEEDHAM

2.1 Introduction

Ideologues may deal with the world as they would wish it to be, but engineers deal with the world as it is. If you're going to defend systems against attack, you first need to know who your enemies are.

In the early days of computing, we mostly didn't have real enemies; while banks and the military took some trouble to protect their systems, most other people didn't bother. The first computer systems were isolated, serving a single company or university. Students might try to hack the system to get more resources and sysadmins would try to stop them, but it was mostly a game. When dial-up connections started to appear, pranksters occasionally guessed passwords and left joke messages, as they'd done at university. The early Internet was a friendly place, inhabited by academics, engineers at tech companies, and a few hobbyists. We knew that malware was possible but almost nobody took it seriously until the late 1980s when PC viruses appeared, followed by the Internet worm in 1989. (Even that was a student experiment that escaped from the lab.)

Things changed once everyone started to get online. The mid-1990s saw the first spam, the late 1990s brought the first distributed denial-of-service attack, and the explosion of mail-order business in the dotcom boom introduced credit card fraud. To begin with, online fraud was a cottage industry; the same person would steal credit card numbers and use them to buy goods which he'd then sell, or make up forged cards to use in a store. Things changed in the mid-2000s with the emergence of underground markets. These let the bad guys specialise

– one gang could write malware, another could harvest bank credentials, and yet others could devise ways of cashing out. This enabled them to get good at their jobs, to scale up and to globalise, just like manufacturing did in the late eighteenth century. The 2000s also saw the world’s governments putting in the effort to ‘master the Internet’ (as the NSA put it) – working out how to collect data at scale and index it, just like Google does, to make it available to analysts. It also saw the emergence of social networks, so that everyone could have a home online – not just geeks with the skills to create their own handcrafted web pages. And of course, once everyone is online, that includes not just the spooks and the crooks but also the jerks, creeps, racists and bullies.

Over the past decade, this threat landscape has stabilised. We also know quite a lot about it. Thanks to Ed Snowden and other whistleblowers, we know a lot about the capabilities and methods of Western intelligence services; we’ve also learned a lot about China, Russia and other nation-state threat actors. We know a lot about cybercrime; online property crime now makes up about half of all crime, by volume and by value. There’s a substantial criminal infrastructure based on malware and botnets with which we are constantly struggling; there’s also a large ecosystem of scams. Many traditional crimes have gone online, and a typical firm has to worry not just about external fraudsters but also about dishonest insiders. Some firms have to worry about hostile governments, some about other firms, and some about activists. Many people have to deal with online hostility, from kids suffering cyber-bullying at school through harassment of elected politicians to people who are stalked by former partners. And our politics may become more polarised because of the dynamics of online extremism.

One of the first things the security engineer needs to do when tackling a new problem is to identify the likely opponents. Although you can design some specific system components (such as cryptography) to resist all reasonable adversaries, the same is much less true for a complex real-world system. You can’t protect it against all possible threats and still expect it to do useful work at a reasonable cost. So what sort of capabilities will the adversaries have, and what motivation? How certain are you of this assessment, and how might it change over the system’s lifetime? In this chapter I will classify online and electronic threats depending on motive. First, I’ll discuss surveillance, intrusion and manipulation done by governments for reasons of state, ranging from cyber-intelligence to cyber-conflict operations. Second, I’ll deal with criminals whose motive is mainly money. Third will be actors whose reasons are personal and who mainly commit crimes against the person, from hacktivists to stalkers and cyber-bullies.

2.2 The Spooks

Governments have a range of tools for both passive surveillance of networks and active attacks on computer systems. Hundreds of firms sell equipment for wiretapping, for radio intercept, and for using various vulnerabilities to take over computers, phones and other digital devices. However there are significant differences between governments in scale, objectives and capabilities. We’ll dis-

cuss four representative categories – the USA and its allies, China, Russia and the Arab world – from the viewpoint of potential opponents. Even if the spooks aren't in your threat model today, the tools they use will quite often end up in the hands of the crooks too, sooner or later.

2.2.1 The Five Eyes

Just as everyone who's old enough remembers where they were when John Lennon was shot, so everyone who's been in our trade since 2013 remembers where they were when they learned of the Snowden revelations on Friday 7th June of that year.

2.2.1.1 Prism

I was in a hotel in Palo Alto, California, reading the Guardian online before a scheduled visit to Google where I'd been as a scientific visitor in 2011, helping develop contactless payments for Android phones. The headline was “NSA Prism program taps in to user data of Apple, Google and others”; the article, written by Glenn Greenwald and Ewen MacAskill, describes a system called Prism that collects the gmail and other data of users who are not US citizens or permanent residents, and is carried out under an order from the FISA court [593]. After breakfast I drove to the Googleplex, and found that my former colleagues were just as perplexed as I was. They knew nothing about Prism. Neither did the gmail team. How could such a wiretap have been built? Had an order been served on Eric Schmidt, and if so how could he have implemented it without the mail and security teams knowing? As the day went on, people stopped talking.

It turned out that Prism was an internal NSA codename for an access channel that had been provided to the FBI to conduct warranted wiretaps. US law permits US citizens to be wiretapped provided an agency convinces a court to issue a warrant, based on ‘probable cause’ that they were up to no good; but foreigners could be wiretapped freely. So for a foreign target like me, all an NSA intelligence analyst had to do is click on a tab saying he believed I was a non-US person. The inquiry would be routed automatically via the FBI infrastructure and pipe my gmail to their workstation. According to the article, this program had started at Microsoft in 2007; Yahoo had fought it in court, but lost, joining in late 2008; Google and Facebook had been added in 2009 and Apple finally in 2012. A system that people thought was providing targeted, warranted wiretaps to law enforcement was providing access at scale for foreign intelligence purposes, and according to a slide deck leaked to the Guardian it was ‘the SIGAD¹ most used in NSA reporting’.

The following day we learned that the source of the story was Edward Snowden, an NSA system administrator who'd decided to blow the whistle. The story was that he'd smuggled over 50,000 classified documents out of a facility in Hawaii on a memory stick and met Guardian journalists in Hong Kong [594]. He tried to fly to Latin America on June 21st to claim asylum, but after the US government cancelled his passport he got stuck in Moscow and eventually got

¹SIGINT (Signals Intelligence) Activity Designator

asylum in Russia instead. A consortium of newspapers coordinated a series of stories describing the signals intelligence capabilities of the ‘Five eyes’ countries – the USA, the UK, Canada, Australia and New Zealand – as well as how these capabilities were not just used but also abused. The first story based on the leaked documents had actually appeared two days before the Prism story; it was about how the FISA court had ordered Verizon to hand over all call data records to the NSA in February that year [590]. This hadn’t got much attention from security professionals as we knew the agencies did that anyway. But it did show that US Director of National Intelligence James Clapper had misled Congress when he’d testified that the NSA collects Americans’ domestic communications ‘only inadvertently’. And what was to follow changed everything.

2.2.1.2 Tempora

One June 21st, the press ran stories about Tempora, a program to collect intelligence from international fibre optic cables [874]. This wasn’t a complete surprise; the journalist Duncan Campbell had described a system called Echelon in 1988 which tapped the Intelsat satellite network, keeping voice calls on tape while making metadata available for searching so that analysts could select traffic to or from phone numbers of interest [281, 282] (I’ll describe more historical background in 24.2.6). Snowden gave us an update on the technology. In Cornwall alone, 200 transatlantic fibres were tapped and 46 could be collected at any one time. As each of these carried 10Gb, the total data volume could be as high as 21Pb a day, so the incoming data feeds undergo ‘massive volume reduction’, discarding video, news and the like. Material was then selected using ‘selectors’ – not just phone numbers but more general search terms such as IP addresses – and stored for 30 days in case it’s of interest. The Tempora program, like Echelon before it, has heavy UK involvement. Britain has physical access to about a quarter of the Internet’s backbone, as modern cables tend to go where phone cables used to, and they were often laid between the same end stations as nineteenth-century telegraph cables. So one of the UK’s major intelligence assets turns out to be a legacy of its nineteenth-century empire. And the asset is indeed significant: by 2012, 300 analysts from GCHQ, and 250 from the NSA, were sifting through the data, using 40,000 and 31,000 selectors respectively to sift 600m ‘telephone events’ each day.

2.2.1.3 Muscular

One of the applications running on top of Tempora was Muscular. Revealed on October 30th, this collected data as it flowed between the data centres of large service firms such as Yahoo and Google [548]. Your mail may have been encrypted using SSL en route to the service’s front end, but it then flowed in the clear between each company’s data centres. After an NSA powerpoint slide on ‘Google Cloud Exploitation’ was published in the Washington Post – see figure 2.1 – the companies scrambled to encrypt everything on their networks. Executives and engineers at cloud service firms took the smiley as a personal affront. It reminded people that even if you comply with warrants, the spooks will also hack you if they can.

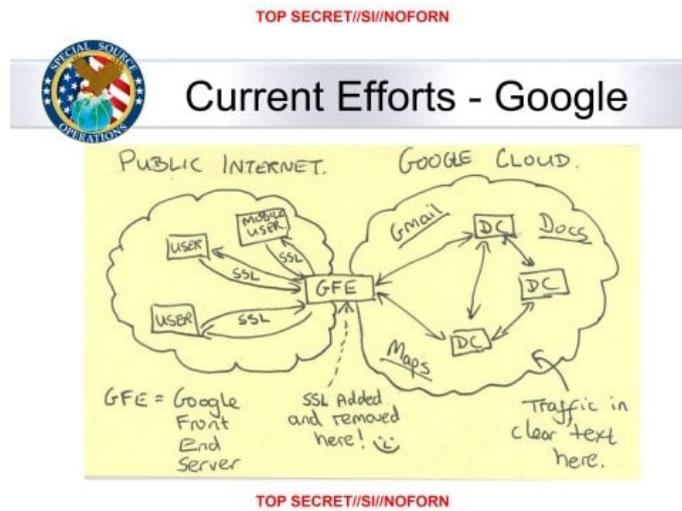


Figure 2.1: Muscular – the slide

Two years later, at a meeting at Princeton which Snowden attended in the form of a telepresence robot, he pointed out that a lot of Internet communications that appear to be encrypted aren't really, as modern websites use content delivery networks (CDNs) such as Akamai and Cloudflare; while the web traffic is encrypted from the user's laptop or phone to the CDN's point of presence at their ISP, it isn't encrypted on the backhaul unless they pay extra – which most of them don't [65]. So the customer thinks the link is encrypted, and it's protected from casual snooping – but not from nation states or firms who can read backbone traffic.

2.2.1.4 Special collection

The NSA and CIA jointly operate the Special Collection Service (SCS) whose most visible activity may be the plastic panels near the roofs of US and allied embassies worldwide; these hide antennas for hoovering up cellular communication (a program known as 'Stateroom'). Beyond this, SCS implants collection equipment in foreign telcos, Internet exchanges and government facilities. This can involve classical spy tradecraft, from placing bugs that monitor speech or electronic communications, through recruiting moles in target organisations, to the covert deployment of antennas in target countries to tap internal microwave links. Such techniques are not restricted to state targets: Mexican drug cartel leader 'El Chapo' Guzman was caught after US agents suborned his system administrator.

Close-access operations include Tempest monitoring: the collection of information leaked by the electromagnetic emissions from computer monitors and other equipment, described in 17.4.2. The Snowden leaks disclose the collection of computer screen data and other electromagnetic emanations from a number of countries' embassies and UN missions including those of India, Japan, Slovakia

and the EU.²

Special collection increasingly involves supply-chain tampering. SCS routinely intercepts equipment such as routers being exported from the USA, adds surveillance implants, repackages them with factory seals and sends them onward to customers. And an extreme form of supply-chain tampering was when the NSA covertly bought Crypto AG, a Swiss firm that was the main supplier of cryptographic equipment to non-aligned countries during the Cold War; I tell the story in more detail later in Section 24.2.7.1.

2.2.1.5 Xkeyscore

With such a vast collection of data, you need good tools to search it. The Five Eyes search computer data using Xkeyscore, a distributed database that enables an analyst to search collected data remotely and assemble the results. Exposed on July 31 2013, NSA documents describe it as its “widest-reaching” system for developing intelligence; it can parse html, and enables an analyst to search emails, SMSes, chats, address book entries and browsing histories [591]. Examples in a 2008 training deck include ‘my target speaks German but is in Pakistan. How can I find him?’ ‘Show me all the encrypted Word documents from Iran’ and ‘Show me all PGP usage in Iran’. By searching for anomalous behaviour, the analyst can find suspects and identify strong selectors (such as email addresses, phone numbers or IP addresses) for more conventional collection.

Xkeyscore is a federated system, where one query scans all sites. Its components buffer information at collection points; in 2008 there were 700 servers at 150 sites. Some appear to be hacked systems overseas from which the NSA malware can exfiltrate data matching a submitted query. The only judicial approval required is a prompt for the analyst to enter a reason why she believes that one of the parties to the conversation is not resident in the USA. The volumes are such that traffic data are kept for 30 days but content for only 3–5 days. Tasked items are extracted and sent on to whoever requested them, and there’s a notification system (Traffichief) for tipping off analysts when their targets do anything of interest. Extraction is based either on fingerprints or plugins – the latter allow analysts to respond quickly with detectors for new challenges like steganography and homebrew encryption.

Xkeyscore can also be used for target discovery: one of the training queries is ‘Show me all the exploitable machines in country X’ (machine fingerprints are compiled by a crawler called Mugshot). For example, it came out in 2015 that GCHQ and the NSA hacked the world’s leading provider of SIM cards, the Franco-Dutch company Gemalto, to compromise the keys needed to intercept (and if need be spoof) the traffic from hundreds of millions of mobile phones [1190]. The hack used Xkeyscore to identify the firm’s sysadmins, who were then phished; agents were also able to compromise billing servers to suppress SMS billing and authentication servers to steal keys; another technique was to harvest keys in transit from Gemalto to mobile service providers. According to an interview with Snowden in 2014, Xkeyscore also lets an analyst

²If the NSA needs to use high-tech collection against you as they can’t simply emplace a software implant in your computer, that may be a compliment!

build a fingerprint of any target's online activity so that they can be followed automatically round the world. The successes of this system are claimed to include the capture of over 300 terrorists; in one case, Al-Qaida's Sheikh Atiyatallah blew his cover by googling himself, his various aliases, an associate and the name of his book [1190].

There's a collection of decks on Xkeyscore with a survey by Morgan Marquis-Boire, Glenn Greenwald and Micah Lee [894]; a careful reading of the decks can be a good starting point for exploring the Snowden hoard³.

2.2.1.6 Longhaul

Bulk key theft and supply-chain tampering are not the only ways to defeat cryptography. The Xkeyscore training deck gives an example: 'Show me all the VPN startups in country X, and give me the data so I can decrypt and discover the users'. VPNs appear to be easily defeated; a decryption service called Longhaul ingests ciphertext and returns plaintext. The detailed description of cryptanalytic techniques is held as Extremely Compartmented Information (ECI) and is not found in the Snowden papers, but some of them talk of recent breakthroughs in cryptanalysis. What might these be?

The leaks do show diligent collection of the protocol messages used to set up VPN encryption, so some cryptographers suggested in 2015 that some variant of the 'Logjam attack' is feasible for a nation-state attacker against the 1024-bit prime used by most VPNs and many TLS connections with Diffie-Hellman key exchange [20]. The algorithm used to attack systems that use discrete log with a prime exponent, the number field sieve, involves a large initial computation for each prime modulus in use but only a short computation for each key generated using that modulus. The Logjam attack also uses a protocol flaw where the server doesn't sign the ciphersuite in use. Others have suggested that a slightly more efficient number field sieve algorithm, about which there had been rumours in the computational number theory community, would yield the same result without the need for a protocol exploit; yet others point to the involvement of cryptologists who worked for NSA contractors during the design of the key-management protocols most commonly used in VPNs.

2.2.1.7 Quantum

As we discussed in the chapter on protocols, there is a long history of attacks on protocols, which can be spoofed, replayed and manipulated in various ways. We discussed the man-in-the-middle (MITM) attack, and even the Mig-in-the-middle variant on IFF. The best documented NSA attack on Internet traffic goes under the codename of Quantum and involves the dynamic exploitation of one of the communication end-points. Thus, to tap an encrypted SSL/TLS session to a webmail provider, the Quantum system fires a 'shot' that exploits the browser. There are various flavours; in 'Quantuminsert', an injected packet redirects the browser to a 'Foxacid' attack server. Other variants attack software updates and the advertising networks whose code runs in mobile phone apps [1416].

³There's also a search engine for the collection at <https://www.edwardsnowden.com>.

2.2.1.8 CNE

Computer and Network Exploitation (CNE) is the generic NSA term for hacking, and it can be used for more than just key theft or TLS session hijacking; it can be used to acquire access to traffic too. Operation Socialist was the GCHQ codename for a hack of Belgian’s main telco Belgacom (now Proximus), in 2010–11. GCHQ attackers used Xkeyscore to identify three key Belgacom technical staff, then used Quantuminsert to take over their PCs when they visited sites like LinkedIn, installed malware on dozens of servers, including authentication servers to leverage further access, billing servers so they could cover their tracks, and the company’s core Cisco routers [533]. This gave them access to large quantities of mobile roaming traffic, as Belgacom provides service to many foreign providers when their subscribers roam in Europe. The idea that one NATO and EU member state would conduct a cyber-attack on the critical infrastructure of another took many by surprise. The attack also gave GCHQ access to the phone system in the European Commission and other European institutions. Given that these institutions make many of the laws for the UK and other member states, this was almost as if a US state governor had got his state troopers to hack AT&T so he could wiretap Congress and the White House.

Belgacom engineers started to suspect something was wrong in 2012, and realised they’d been hacked in the spring of 2013; an anti-virus company found sophisticated malware masquerading as Windows files. The story went public in September 2013, and Der Spiegel published Snowden documents showing that GCHQ was responsible. After the Belgian prosecutor reported in February 2018, we learned that the attack must have been authorised by then UK Foreign Secretary William Hague, but there was not enough evidence to prosecute anyone; the investigation had been hampered in all sorts of ways both technical and political; the software started deleting itself within minutes of discovery, and institutions such as Europol (whose head was British) refused to help. The Belgian minister responsible for telecomms, Alexander de Croo, even suggested that Belgium’s own intelligence service might have informally given the operation a green light [534]. Europol later adopted a policy that it will help investigate hacks of ‘suspected criminal origin’; it has nothing to say about hacks by governments.

A GCHQ slide deck on CNE explains that it’s used to support conventional sigint both by redirecting traffic and by ‘enabling’ (breaking) cryptography; that it must always be ‘UK deniable’; and that it can also be used for ‘effects’, such as degrading communications or ‘changing users’ passwords on extremist website’ [534].

The hacking tools used by the NSA and its allies are now fairly well understood. The Snowden papers reveal an internal store where analysts can get a variety of tools; a series of leaks in 2016–7 by the Shadow Brokers (thought to be Russian military intelligence, the GRU) disclosed a number of actual NSA malware samples, used by hackers of the NSA’s Tailored Access Operations team to launch attacks [181]. (Some of these tools were repurposed by the Russians to launch the NotPetya worm and by the North Koreans in Wannacry, as I’ll discuss later.) The best documentation of all is probably about a separate store of goodies used by the CIA, disclosed in some detail to wikileaks in the ‘Vault 7’

leaks in 2017. These tools can be used to install a remote access Trojan on your machine, with components to geolocate it and to exfiltrate files (including SSH credentials), audio and video. There's also a tool to jump air gaps by infecting thumb drives. Many of the tools are available not just for Windows but also for OSX and Android; some infect firmware, making them hard to remove. There's also a tool for infecting wifi routers so they'll do man-in-the-middle attacks, and even a tool for watermarking documents so a whistleblower who leaks them could be tracked. There are tools for hacking TVs and IoT devices, and tools to hamper forensic investigations. The Vault 7 documents are useful reading if you'd like to see what the specifications and manuals for modern government malware look like [1431].

2.2.1.9 The analyst's viewpoint

The intelligence analyst thus has a big bag of tools. If she's trying to find the key people in an organisation – whether the policymakers advising on a critical decision, or the lawyers involved in laundering an oligarch's profits – she can use the traffic data in Xkeyscore and Boundless Informant to map contact networks. There are various neat tools to help, such as 'cotraveler' which flags up mobile phones that have traveled together. Once she moves from the hunting phase to the gathering phase, she can use Prism to look at their accounts at Google and Microsoft, while Xkeyscore will let her see what websites they visit, and more: despite the growing use of encryption, the communications to and from a home reveal what app or device is used when and for how long⁴. The agencies are pushing for access to end-to-end messaging systems such as WhatsApp; in countries like the UK, Australia and China, legislators have already authorised this, though it's not at all clear which US companies might comply (I'll discuss policy in the chapter on Surveillance and Privacy).

But once the analyst knows her target, there's a big bag of tools she can install on their laptop or cellphone directly. She can locate it physically, turn it into a room bug and even use it as a remote camera. Once she has confirmed that the target is of interest, she can download the address book and contact history and feed that into Xkeyscore to search recursively for its direct and indirect contacts. Meanwhile she can bug messaging apps, beating the end-to-end encryption by collecting the call contents once they've been decrypted. She can set up an alarm to notify her whenever the target sends or receives messages of interest, or changes location. The coverage is pretty complete. And when it's time for the kill, the target's phone can be used to guide a bomb or a missile. Little wonder Ed Snowden insisted that journalists interviewing him put their phones in the fridge!

Finally, the analyst has also a proxy through which she can access the Internet surreptitiously – typically a machine on a botnet. It might even be the PC in your home office.

⁴See for example Hill and Mattu who wiretapped a modern smart home to measure this [646].

2.2.1.10 Offensive operations

The Director NSA also heads the US Cyber Command, which since 2009 has been one of ten unified commands of the United States' Department of Defense. It is responsible for offensive cyber operations, of which the one that made a real difference was Stuxnet. This was a worm designed to damage Iran's uranium enrichment centrifuges by speeding them up and slowing them down in patterns designed to cause mechanical damage, and appears to have been developed jointly by the USA and Israel [251]. It was technically sophisticated, using four zero-day exploits and two stolen code-signing certificates to spread promiscuously through Windows PCs, until it found Siemens programmable logic controllers of the type used at Iran's Natanz enrichment plant – where it would then install a rootkit that would issue the destructive commands, while the PC assured the operators that everything was fine. It was apparently introduced using USB drives to bridge the air gap to the Iranian systems, and came to light in 2010 after copies had somehow spread to central Asia and Indonesia. Two other varieties of malware (Flame and Duqu) were then discovered using similar tricks and common code, performing surveillance at a number of companies in the Middle East and South Asia; more recent code-analysis tools have traced a lineage of malware that goes back to 2002 (Flowershop) and continued to operate until 2016 (with the Equation Group tools) [1466]. Presumably they have newer toolkits now.

Stuxnet acted as a wake-up call for other governments, which rushed to acquire 'cyber-weapons' and develop offensive cyber doctrine. The price of zero-day vulnerabilities rose sharply.

2.2.1.11 Attack scaling

Computer scientists learn the importance of how algorithms scale, and exactly the same holds for attacks. Tapping a single mobile phone is hard. You have to drive around behind the suspect with radio and cryptanalysis gear in your car, risk being spotted, and hope that you manage to catch the suspect's signal as he roams from one cell to another. Or you can use a MITM attack like a Stingray and risk electronic detection too. Both are highly skilled work and low-yield: you lose the signal maybe a quarter of the time. So if you want to wiretap someone in central Paris often enough, why not just wiretap everyone? Put antennas on your embassy roof, collect it all, write the plaintext into a database, and reconstruct the sessions electronically. If you want to hack everyone in France, hack the telco. At each stage the capital cost goes up but the marginal cost of each tap goes down. The Five Eyes strategy is essentially to collect everything in the world; that might cost billions to establish and maintain the infrastructure, but once it's there you have everything.

The same applies to offensive cyber operations, which are rather like sabotage. In wartime, you can send commandos to blow up an enemy radar station; but if you do it more than once or twice, your lads will start to run into a lot of sentries. So we scale kinetic attacks differently: by building hundreds of bomber aircraft, or artillery pieces, or (nowadays) thousands of drones. So how do you scale a cyber attack to take down not just one power station, but the

opponent's whole power grid? The Five Eyes approach is this. Just as Google keeps a copy of the Internet on a few thousand servers, with all the content and links indexed, so US Cyber Command keeps a copy of the Internet that indexes what version of software all the machines in the world are using – the Mugshot system mentioned above – so a Five Eyes cyber warrior can instantly see which targets can be taken over by which exploits.

A key question for competitor states, therefore, is not just to what extent they can create some electronic spaces that are generally off-limits to the Five Eyes; it's the extent to which they can scale up their own intelligence and offensive capabilities rather than having to rely on America. The number of scans and probes that we see online indicates that the NSA are not alone in trying to build cyber weapons that scale. Not all of them might be nation states; some might simply be arms vendors or mercenaries. This raises a host of policy problems to which we'll return in part 3. For now we'll continue to look at capabilities.

2.2.2 China

China is now the leading competitor to the USA, being second not just in terms of GDP but as a technology powerhouse. The Chinese lack the NSA's network of alliances and access to global infrastructure (although they're working at that). Within China itself, however, they demand unrestricted access to local data. Some US service firms used to operate there, but trouble followed. After Yahoo's systems were used to trap the dissident Wang Xiaoning in 2002, Alibaba took over Yahoo's China operation in 2005; but there was still a row when Wang's wife sued Yahoo in US courts in 2007, and showed that Yahoo had misled Congress over the matter [1255]. In 2008, it emerged that the version of Skype available in China had been modified so that messages were scanned for sensitive keywords and, if they were found, the user's texts were uploaded to a server in China [1386]. In December 2009, Google discovered a Chinese attack on its corporate infrastructure, which became known as Operation Aurora; Chinese agents had hacked into the systems used to do wiretaps (see Prism above) in order to discover which of their own agents in the USA were under surveillance. Google had already suffered criticism for operating a censored version of their search engine for Chinese users, and a few months later, they pulled out of China. By this time, Facebook, Twitter and YouTube had already been blocked. A Chinese strategy was emerging of total domestic control, augmented by ever-more aggressive collection overseas.

From about 2002, there had been a series of hacking attacks on US and UK defence agencies and contractors, codenamed 'Titan Rain' and ascribed to the Chinese armed forces [1040]. According to a 2004 study by the US Foreign Military Studies Office (FMSO), Chinese doctrine sees the country in a state of war with the west; in an extension of the Cold War, the West is already attacking China by exporting subversive ideas to it over the Internet [1331]. Chinese leaders see US service firms, news websites, anonymity tools such as Tor⁵ as being of one fabric with the US surveillance satellites and aircraft that observe their military defences. Yahoo and Google were thus seen as fair game,

⁵which the State Department funds so that Chinese and other people can defeat censorship

just like Lockheed Martin and BAe.

Our own first contact with Chinese came in 2008. We were asked for help by the Dalai Lama who had realised that the Chinese had hacked his office systems in the run-up to the Beijing Olympics that year. One of my research students, Shishir Nagaraja, happened to be in Delhi waiting for his UK visa to be renewed, and volunteered to go up to the Tibetan HQ in Dharamsala and run some forensics. He found that about 35 of the 50 PCs in the office of the Tibetan government in exile had been hacked; information was being siphoned off to China, to IP addresses located near the three organs of Chinese state security charged with different aspects of Tibetan affairs. The attackers appear to have spear-phished their way in, and then compromised the Tibetans' mail server, so that whenever one person in the office sent a pdf file to another, it arrived with a javascript buffer overflow that would use a vulnerability in Adobe Reader to take over the recipient's machine. The mail server itself was in California.

This is pretty sobering, when you stop to think about it. You get an email from a colleague sitting ten feet away, you ask him if he just sent it – and when he says yes, you click on the attachment. And your machine is suddenly infected by a server that you rent and that's ten thousand miles away in a friendly country. We wrote this up in a tech report on the 'Snooping Dragon' [986]. After it came out, we had to deal for a while with attacks on our equipment, and heckling at conference talks by Chinese people who claimed we had no evidence to attribute the attacks to their government. Colleagues at the Open Net Initiative in Toronto followed through and eventually found from analysis of the hacking tools' dashboard that the same espionage network had targeted 1,295 computers in 103 countries [892] – ranging from the Indian embassy in Washington through Associated Press in New York to the ministries of foreign affairs in Thailand, Iran and Laos.

There followed a series of further reports of Chinese state hacking, from a complex dispute with Rio Tinto in 2009 over the price of iron ore and a hack of the Melbourne International Film festival in the same year when it showed a film about a Uighur leader [1344]. One flashbulb moment was a leaked Pentagon report in 2013 that Chinese hackers had stolen some of the secrets of the F35 joint strike fighter, as well as a series of other weapon systems [990]. Meanwhile China and Hong Kong were amounting for over 80% of all counterfeit goods seized at US ports. The Obama administration vowed to make investigations and prosecutions in the theft of trade secrets a top priority, and the following year five members of the People's Liberation Army were indicted in absentia.

The White House felt compelled to act once more after the June 2015 news that the Chinese had hacked the Office of Personnel Management (OPM), getting access to highly personal data on 22 million current and former federal employees, ranging from fingerprints to sensitive information from clearance. Staff applying for Top Secret clearances are ordered to divulge all information that could be used to blackmail them, from teenage drug use to closeted gay relationships. All sexual partners in the past five years have to be declared for a normal Top Secret clearance; for a Strap clearance (to deal with signals intelligence material) the candidate even has to report any foreigners they meet regularly at their church. So the information leak affected more than just 22 mil-

lion people. Officially, this is to mitigate the risk that intelligence agency staff can be blackmailed. (Cynics supposed it was also so that whistleblowers could be discredited.) Whatever the motives, putting all such information in one place was beyond stupid. For the Chinese to get all the compromising information on every American with a sensitive government job was jaw-dropping. (The UK screwed up it too; in 2008, a navy officer lost a laptop containing the personal data of 600,000 people who had joined the Royal Navy, or tried to. [780].) At a summit in September that year, Presidents Obama and Xi agreed to refrain from computer-enabled theft of intellectual property for commercial gain. Nothing was said in public though about military secrets – or the sex lives of federal agents for that matter.

The Chinese attacks of the 2000s used smart people plus simple tools; the attacks on the Tibetans used Russian crimeware as the remote access Trojans. The state also co-opted groups of ‘patriotic hackers’, or perhaps used them for deniability; some analysts noted waves of naïve attacks on western firms that were correlated with Chinese university terms, and wondered whether students had been tasked to hack as coursework. The UK police and security service warned UK firms in 2007. By 2009, multiple Chinese probes had been reported on US electricity firms, and by 2010, Chinese spear-phishing attacks had been reported on government targets in the USA, Poland and Belgium [941]. As with the Tibetan attacks, these typically used crude tools and had such poor operational security that it was fairly clear where they came from. By 2019 the attacks had become more sophisticated, with a series of advanced persistent threats (APTs) tracked by threat intelligence firms. One of them has been used to exploit managed service providers as an attack channel; this modus operandi became public in 2019 with the disclosure that someone had hacked Wipro and used this to compromise their customers [795]. (Of the big Indian IT service firms, Wipro was the one that had tried hardest to build a reputation for security.) Another approach was attacking software supply chains; a Chinese group variously called Wicked Panda or Barium compromised software updates from computer maker Asus, a PC cleanup tool and a Korean remote management tool, as well as three popular computer games, getting its malware installed on millions of machines; rather than launching banking trojans or ransomware, it was then used for spying [586]. Just as in GCHQ’s Operation Socialist, these indirect strategies give a way to scale attacks in territory where you’re not the sovereign.

During 2018–9 there was a political row over whether Chinese firms should be permitted to sell routers and 5g network hardware in NATO countries, with the Trump administration warning of a back door that could be used later to wiretap the West at scale. In 2018 GCHQ warned that ZTE equipment ‘would present risk to UK national security that could not be mitigated effectively or practicably’ [1059]⁶. GCHQ also set up a centre in 2009 to study Huawei’s software in detail, as a condition of their being allowed to sell in the UK; while they did not find any back doors, their 2019 report surfaced some scathing criticisms of the company’s software engineering practices [671]. No progress

⁶The only router vendor to have actually been caught with a malicious backdoor in its code is the US company Juniper, which not only used the NSA’s Dual-EC backdoor to make VPN traffic exploitable, but did it in such a clumsy way that others could exploit it too – and at least one other party did so [309].

was being made in tackling a wide range of security-engineering problems despite promises made over a number of years. Problems included the inclusion of an unmanageable number of versions of Open SSL, including versions that are not on the main development train, that had known vulnerabilities and that were not supported. In particular there were 70 full copies of 4 different OpenSSL versions, and 304 partial copies of 14 versions. Huawei had copied a lot of code, and couldn't patch what they didn't understand. Not only could the Chinese hack the Huawei systems; so could anybody. This equipment has been excluded for some years from use in UK backbone routers and in wiretapping. The UK wants 'sustained evidence of improvement across multiple versions and multiple product ranges' before it will put any more trust in it. A number of countries, including Australia and New Zealand, have banned it outright, and in 2019 Canada arrested Huawei's CFO (who is also the founder's daughter) following a US request to extradite her for breaking sanctions against Iran. China retaliated by sentencing a Canadian citizen to death. Meanwhile, China is helping less developed countries modernise their networks, which will presumably enable them to rival the Five Eyes' scope in due course. Trade policy, industrial policy and cyber-defence strategy have become intertwined.

Strategically, the question may not be whether China could use this equipment to wiretap the west at scale, so much as whether they could use it in time of tension to launch DDoS attacks that would break the Internet by subverting BGP routing. I discuss this in more detail in the chapter on Network Security. For the meantime, China's doctrine of 'Peaceful Rise' means they avoid conflict with other major powers until they're strong enough. The overall picture is therefore of defensive information warfare, combining pervasive surveillance at home, a walled-garden domestic Internet that is better defended against cyber-attack than anyone else's, plus considerable and growing capabilities – which are not yet used for disruption against major players, but for diligent intelligence-gathering in support of national strategic interests. This defensive posture may of course change; it did in the past with the start of the Korean war, and might do again if there's another regional conflict.

2.2.3 Russia

Russia, like China, lacks America's platform advantage and compensates with hacking teams that use spear-phishing and malware. Its strategic posture is different in three ways. First, it's a major centre for cybercrime; underground markets first emerged in Russia and the Ukraine in 2003–5, and we'll discuss this more in the following section on cybercrime. Second, although Russia is trying to become more closed like China, its domestic Internet including major service firms such as VK and Yandex are relatively open and intertwined with the West's [438]. Third, Russia's strategy of re-establishing itself as a regional power has been pursued much more aggressively than China's, with direct military interference in neighbours such as Georgia and the Ukraine. These interventions have involved a mixed strategy of cyber-attacks plus 'little green men' – troops without Russian insignia on their uniforms – with a political strategy of denial.

Russian cyber-attacks came to prominence in 2007, after Estonia moved a much-hated Soviet-era statue in Tallinn to a less prominent site, and the

Russians felt insulted. DDoS attacks on government offices, banks and media companies forced Estonia to rate-limit its external Internet access for a few weeks [497]; Russia refused to extradite the perpetrators, most of whom were Russian, though one ethnic-Russian Estonian teenager was fined. Sceptics said that the attacks seemed the work of amateurs and worked because the Estonians hadn't hardened their systems in the way that US service providers do. Estonia nonetheless appealed to NATO for help, and one outcome was the Tallinn Manual, which sets out the military and international law applicable to online operations designed to have real-world effects in conflicts between states [1198]. States must act in self-defence or with some other lawful justification and in accordance with the law of armed conflict⁷. The following year, after the outbreak of a brief war between Russia and Georgia, Russian hackers set up a website with a list of targets in Georgia for patriots to attack [1412].

Estonia and Georgia were little more than warm-ups for the Ukraine. Following the Russian annexation of Crimea in 2014 and the imposition of international sanctions, Russia supported an insurgency in the Donbass area of eastern Ukraine. The following year, cyber-attacks took down 30 electricity substations, leaving 230,000 people without electricity for several hours. This was tiny compared with the other effects of the conflict – which included the shooting down of a Malaysian Airlines airliner with the loss of all on board – but it may be significant as the first cyber-attack to disrupt mains electricity. Finally on June 27 2017 came the NotPetya attack – by far the most damaging cyber attack to date [589].

The NotPetya worm was initially distributed using the update service for MeDoc, the accounting software used by the great majority of Ukrainian businesses. It then spread laterally in organisations across Windows file-shares using the EternalBlue vulnerability, an NSA exploit with an interesting history. From March 2016, a Chinese gang started using it against targets in Vietnam, Hong Kong and the Philippines, perhaps as a result of finding and reverse engineering it (it's said that you don't launch a cyberweapon; you share it). It was then leaked by a gang called the 'Shadow Brokers' in April 2017, along with other NSA software that the Chinese didn't deploy, and finally used by the Russians in June. The NotPetya worm used EternalBlue together with an exploit called Mimikatz that recovers passwords from Windows memory. The worm's payload pretended to be ransomware; it encrypted the infected computer's hard disk and demanded a ransom of \$300 in bitcoin. But there was no mechanism to decrypt the files of computer owners who paid the ransom, so it was really a destructive service-denial worm. The only way to deal with it was to re-install the operating system and restore files from backup.

The NotPetya attack took down banks, telcos and even the radiation monitoring systems at the former Chernobyl nuclear plant. What's more, it spread from the Ukraine to international firms who had offices there. The shipping company Maersk had to replace most of its computers and compensate cus-

⁷Attacks are operations reasonably expected to cause injury to people or damage to property; they may only be directed at combatants and their logistics, not at civilians; attacks must be geographically limited, not indiscriminate; and some targets are off-limits, from hospitals and places of worship to nuclear power stations. Interpretation could keep the lawyers busy though. Infrastructure used by both military and civilian organisations is fair game, and although 'treachery' is prohibited, 'ruses of war' are not.

tomers for late shipments, at a cost of \$300m; other firms affected included FedEx (which lost \$300m) and Mondelez (\$100m) – whose insurers refused to pay out on the ground that it was an ‘Act of War’. Their justification was that the governments of the Ukraine, the USA and the UK attributed NotPetya to Russian military intelligence, the GRU [895].

2016 was marked by the Brexit referendum in the UK and the election of President Trump in the USA, in both of which there was some Russian interference. In the former, the main intervention appears to have been financial support for the leave campaigns, which were later found to have broken the law by spending too much. In the latter, Russian interference was denounced by President Obama during the campaign, leading to renewed economic sanctions, and by the US intelligence community afterwards. An inquiry by former FBI director Robert Mueller found that Russia interfered very widely via the disinformation and social media campaigns run by its Internet Research Agency ‘troll farm’, and by the GRU which hacked the emails of the Democratic national and campaign committees, most notably those of the Clinton campaign chair John Podesta. Some Trump associates went to jail; Mueller found insufficient evidence that Trump conspired with the Russians (as opposed to colluding, which is legal) However he did not exonerate Trump either. As I’ll discuss in the chapter ‘Surveillance or Privacy?’, it is hard to assess the effects of such interventions. On the one hand Yochai Benkler cautions Democrats against believing that Trump’s election was won for him by Russia; the roots of popular disaffection with the political elite are much older and deeper [175]. On the other hand, Carole Cadwalladr has documented how social media provided a channel for Russia and its allies to influence elections and referenda using targeted ads, trolling and other techniques that flout legal limits on campaign contributions; she asks whether it’s even possible to have free and fair elections with unlimited, unmonitored targeted ads [273]. In this turbulent time, Russian state policy is to undermine the influence of democratic states and the rules-based international order, promoting authoritarian governments of both left and right, and causing trouble where it can. In short, its information war with the West is aggressive, and continues the old USSR’s strategy of weakening the West by fomenting conflict by a variety of national liberation movements and terrorist groups.

2.2.4 The rest

The rest of the world’s governments have quite a range of cyber capabilities, but common themes, including the nature and source of their tools. Middle Eastern governments were badly shaken by the Arab Spring uprisings, and some even turned off the Internet for a while, such as Libya in April–July 2010, when rebels were using Google maps to generate target files for US, UK and French warplanes. Since then, Arab states have developed strategies that combine spyware and hacking against high-profile targets, through troll farms pumping out abuse comments in public fora, to simple arrest.

The operations of the United Arab Emirates were described in 2019 by a whistleblower, Lori Stroud [186]. An NSA analyst – and Ed Snowden’s former boss – she was headhunted by a Maryland contractor in 2014 to work in Dubai

as a mercenary, but left after the UAE's operations started to target Americans. The UAE's main technique was spear-phishing with Windows malware, but their most effective tool, called Karma, enabled them to hack the iPhones of foreign statesmen and local dissidents. They also targeted westerners critical of the regime, in one case social-engineering a UK grad student into installing spyware on his PC on the pretext that it would make his communications hard to trace. The intelligence team consisted of several dozen people, both mercenaries and Emiratis, in a large villa in Dubai. The use of iPhone malware by the UAE government was documented by independent observers [890].

In 2018, the government of Saudi Arabia was exposed for murdering the Washington Post journalist Jamal Khashoggi in its consulate in Istanbul. The Post campaigned to expose Saudi crown prince Mohammed bin Salman as the man who gave the order, and in January 2019 the National Enquirer published a special edition containing texts showing that the Post's owner Jeff Bezos was having an affair. He announced a separation from his wife, which led to the most expensive divorce in history, and hired an investigator to find the source of the leak. The National Enquirer then attempted to blackmail Bezos over some photos it has also obtained; it wanted both him and the investigator to declare that the paper hadn't relied upon 'any form of electronic eavesdropping or hacking in their news-gathering process'. Bezos went public instead; according to the investigator, his iPhone had been hacked by a Saudi Arabian government agency [150].

An even more unpleasant example is Syria, where the industrialisation of brutality is a third approach to scaling information collection. Malware attacks on dissidents were reported from 2012, and initially used a variety of spear-phishing lures. As the civil war got underway, police who were arresting suspects would threaten female family members with rape on the spot unless the suspect disclosed his passwords for mail and social media. They would then spear-phish all his contacts while he was being taken away in the van to the torture chamber. This victim-based approach to attack scaling resulted in the compromise of many machines not just in Syria but in America and Europe. The campaigns became steadily more sophisticated as the war evolved, with false-flag attacks, yet retained a brutal edge with some tools displaying beheading videos [535].

Thanks to John Scott-Railton and colleagues at Toronto, we have many further documented examples of online surveillance, computer malware and phone exploits being used to target dissidents; many in Middle Eastern and African countries but also in Mexico and indeed in Hungary [890]. The real issue here is the ecosystem of companies, mostly in the USA, Europe and Israel, that supply hacking tools to unsavoury states. NGOs have made attempts to push back on this cyber arms trade; in one case we argued that the Syrian government's ability to purchase mass-surveillance equipment from the German subsidiary of a UK company should be subject to export control, but the UK authorities were unwilling to block it. But GCHQ was determined that if there were going to be bulk surveillance devices on President Assad's network, they should be British devices rather than Ukrainian ones. (I describe this later in the chapter 'Surveillance or Privacy?') So the ethical issues around conventional arms sales persist in the age of cyber; indeed they become worse because these tools can be and are used against Americans, Brits and others who are sitting at home

but who are unlucky enough to be on the contact list of someone an unpleasant government doesn't like. Selling physical weapons to a far-off dictator didn't used to put people in your own country in harm's way; but the same does not hold true for cyber weapons.

Finally, it's worth mentioning North Korea. In 2014, after Sony Pictures started working on a comedy about a plot to assassinate the North Korean leader, a hacker group trashed much of Sony's infrastructure, released embarrassing emails that caused its top film executive Amy Pascal to resign, and leaked some unreleased films. This was followed by threats of terrorist attacks on movie theatres if the comedy were put on general release. The company put the film on limited release instead, and President Obama criticised them for giving in to North Korean blackmail.

In 2017, North Korea again came to attention after their Wannacry worm infected over 200,000 computers worldwide, encrypting data and demanding a bitcoin ransom (though like NotPetya it didn't have a means of selective decryption, so was really just a destructive worm). It used the NSA EternalBlue vulnerability which had been released by the Shadow Brokers, and was stopped when an antivirus researcher discovered a kill switch. In the meantime it had disrupted production at carmakers Nissan and Renault and at the Taiwanese chip foundry TSMC, and also caused several hospitals in Britain's National Health Service to close their accident and emergency units. In 2018, the US Department of Justice unsealed an indictment of a North Korean government hacker for both incidents, and also for a theft of \$81m from the Bank of Bangladesh [1188].

2.2.5 Attribution

It's often said that cyber is different, because attribution is hard. As a general proposition this is untrue; anonymity online is much harder than you think. Even smart people make mistakes in operational security that give them away. Yet sometimes it may be true, and people still point to the Climategate affair. Several weeks before the 2009 Copenhagen summit on climate change, someone published over a thousand emails, mostly sent to or from four climate scientists at the University of East Anglia, England. Climate sceptics seized on some of them, which discussed how to best present evidence of global warming, as evidence of a global conspiracy. Official inquiries later established that the emails had been quoted out of context, but the damage had been done. People wonder whether the perpetrator could have been the Russians or the Saudis or even an energy company. However one of the more convincing analyses suggests that it was an internal leak, or even an accident; only one archive file was leaked, and its filename (FOIA2009.zip) suggests it may have been prepared for a freedom-of-information disclosure in any case. The really interesting thing here may be how the emails were talked up into a conspiracy theory.

Another possible state action was the Equifax hack. On March 8th 2017, Apache warned of a vulnerability in Apache Struts and issued a patch; two days later, a gang started looking for vulnerable systems; on May 13th, they found that Equifax's dispute portal had not been patched, and got in. They found a plaintext password file giving access to 51 internal database systems, and spend 76 days helping themselves to the personal information of at least

145.5 million Americans before the intrusion was reported on July 29 and access blocked the following day. Executives sold stock before they notified the public on September 7th; so Congress was outraged, and the CEO Rick Smith was fired [1038]. So far, so ordinary. But the interesting this is that no criminal use has been made of any of the stolen information, which has led some analysts to suspect that the crime was committed by (or the data sold to) a nation-state actor seeking personal data on Americans at scale, perhaps as a basis for intelligence activities.

In any case, the worlds of intelligence and crime have long been entangled, and in the cyber age they seem to be getting more so. We turn to cybercrime next.

2.3 Crooks

Cybercrime is now about half of all crime, both by volume and by value, at least in those developed countries that ask about fraud in regular victimisation surveys⁸. Whether it is slightly more or less than half depends on definitions (do you include tax fraud now that tax returns are filed online?) and on the questions you ask (do you count harassment and cyber-bullying?) – but even with narrow definitions, it’s still almost half. Yet the world’s law-enforcement agencies typically spend less than one percent of their budgets on fighting it. Much more needs to be done, so colleagues and I have a cybercrime centre at Cambridge where we collect and curate data for other researchers to use, and we also write survey papers from time to time. This section draws on our 2019 analysis [67]; there is also a growing literature on cyber-criminology.

Computer fraud has been around since the 1960s, a notable early case being the Equity Funding insurance company which from 1964-72 created more than 60,000 bogus policies which it sold to reinsurers, creating a special computer system to keep track of them all. Electronic frauds against payment systems have been around since the 1980s, and spam arrived when the Internet was opened to all in the 1990s. Yet early scams were mostly a cottage industry, where individuals or small groups collected credit card numbers, forged cards or used card numbers to get mail-order goods, and figured out how to cash out. Modern cybercrime can probably be dated to 2003–5 when underground markets emerged that enabled crooks to specialise and get good at their jobs, just as happened in the real economy with the Industrial Revolution.

To make sense of cybercrime, it’s convenient to consider the shared infrastructure first, and then the main types of cybercrime that are conducted for profit. There is a significant overlap with the crimes committed by states that we considered in the last section, and those committed by individuals against other individuals that we’ll consider in the next one; but the actors’ motives are a useful primary filter.

⁸The USA, the UK, Australia, Belgium and France

2.3.1 Criminal infrastructure

Since about 2005, the emergence of underground markets has led to people specialising as providers of criminal infrastructure, most notably botnet herders, malware writers, spam senders and cashout operators. I will discuss the technology behind the first three in much greater detail in the chapter on Network Attack and Defense; in this section my focus is on the actors and the ecosystem in which they operate. Although this ecosystem consists of perhaps a few thousand people with revenues in the tens to low hundreds of millions, they impose costs of many billions on the industry and on society.

2.3.1.1 Botnet herders

The first botnets – networks of compromised computers – may have been seen in 1996 with an attack on the ISP Panix in New York, using compromised Unix machines in hospitals to conduct a SYN flood attack [275]. The next use was spam, and by 2000 the Earthlink spammer sent over a million phishing emails; its author was sued by Earthlink. Once cyber-criminals started to get organised, there was a significant scale-up. We started to see professionally built and maintained botnets that could be rented out by bad guys, whether spammers, phishermen or others; by 2007 the Cutwail botnet was sending over 50 million spams a minute from over a million infected machines [1306]. Bots would initially contact a command-and-control server for instructions; these would be taken down, or taken over by threat intelligence companies for use as sinkholes to monitor infected machines, and to feed lists of them to ISPs and corporates.

The spammers' first response was peer-to-peer botnets. In 2007 Storm suddenly grew to account for 8% of all Windows malware; it infected machines mostly by malware in email attachments and had them use the eDonkey peer-to-peer network to find other infected machines. It was used not just for spam but for DDoS, for pump-and-dump stock scams and for harvesting bank credentials. Defenders got lots of peers to join this network to harvest lists of bot addresses, so the bots could be cleaned up, and by late 2008 Storm had been cut to a tenth of the size. It was followed by Kelihos, a similar botnet that also stole bitcoins; its creator, a Russian national, was arrested while on holiday in Spain in 2017 and extradited to the USA where he pled guilty in 2018 [476].

The next criminal innovation arrived with the Conficker botnet: the domain generation algorithm (DGA). Conficker was a worm that spread by exploiting a Windows network service vulnerability; it generated 250 domain names every day, and infected machines would try them all out in the hope that the botmaster had managed to rent one of them. Defenders started out by simply buying up the domains, but a later variant generated 50,000 domains a day and an industry working group made agreements with registrars that these domains would simply be put beyond use. By 2009 Conficker had grown so large, with maybe ten million machines, that it was felt to pose a threat to the largest websites and perhaps even to nation states. As with Storm, its use of randomisation proved to be a two-edged sword; defenders could sit on a subset of the domains and harvest feeds of infected machines. By 2015 the number of infected machines

had fallen to under a million.

Regardless of whether something can be done to take out the command and control system, whether by arresting the botmaster or by technical tricks, the universal fix for botnet infections is to clean up infected machines. But this raises many issues of scale and incentives. While AV companies made tools available, and Microsoft supplies patches, many people don't use them. So long as your infected PC is merely sending occasional spam but works well enough otherwise, why should you go to the trouble of doing anything? But bandwidth costs ISPs money, so the next step was that some ISPs, particularly the cable companies like Comcast, would identify infected machines and confine their users to a 'walled garden' until they promised to clean up. By 2019 that has become less common as people now have all sorts of devices on their wifi, many of which have no user interface; communicating with human users has become harder.

In 2019, we find many botnets with a few tens of thousands of machines that are too small for most defenders to care about, plus some large ones that tend to be multilayer – typically with peer-to-peer mechanisms at the bottom that enable the footsoldier bots to communicate with a few NCO nodes, which in turn use a domain generation algorithm to find the botmaster. Fragmenting the footsoldiers into a number of small clouds makes it hard for defenders to infiltrate all of them, and the NCOs are typically located in places that are hard for defenders to get at. The big money for such botnets appears to be in clickfraud.

The latest innovation – since October 2016 – is Mirai, a family of botnets that exploit IoT devices. The first Mirai worm infected CCTV cameras that had been manufactured by Xiaomai and that had a known factory default password that couldn't be changed. Mirai botnets scan the Internet's IPv4 address space for other vulnerable devices which typically get infected within minutes of being powered up. The first attack was on DynDNS and took down Twitter for six hours on the US eastern seaboard. Since then there have been over a thousand variants, which researchers study to determine what's changed and to work out what countermeasures might be used.

At any one time, there may be half a dozen large botnet herders. The Mirai operators, for example, seem to be two or three groups that might have involved a few dozen people.

2.3.1.2 Malware writers

In addition to the several hundred software engineers who write malware for the world's intelligence agencies and their contractors, there may be from the high hundreds to low thousands of people writing malware for the criminal market; nobody really knows (though we can monitor traffic on hacker forums to guess the order of magnitude).

Within this community there are specialists. Some concentrate on turning vulnerabilities into exploits, a nontrivial task for modern operating systems that use stack canaries, ASLR and other techniques we'll discuss later; others specialise in the remote access Trojans that the exploits install; others build

the peer-to-peer and DGA software for resilient command-and-control communications; others design specialised payloads for bank fraud. The highest-value operations seem to be platforms that are maintained with constant upgrades to cope with the latest countermeasures from the anti-virus companies. Within each specialist market segment there are typically a handful of operators, so that when we arrest one of them it makes a difference for a while. Some of the providers are based in jurisdictions that don't extradite their nationals, like Russia, and Russian crimeware is used not just by Russian state actors but by others too.

As Android has taken over from Windows as the most frequently used operating system we've seen a rise in Android malware. In China and in countries with a lot of second-hand and older phones, this may be software that uses an unpatched vulnerability to root an Android phone; the USA and Europe have plenty unpatched phones (as many OEMs stop offering patches once a phone is no longer on sale) but it's often just apps that do bad things, such as stealing SMSes used to authenticate banking transactions.

2.3.1.3 Bulk account compromise

Some botnets are constantly trying to break into email and other online accounts by trying to guess passwords and password recovery questions. A large email service provider might be recovering several tens of thousands of accounts every day. There are occasional peaks, typically when hackers compromise millions of email addresses and passwords at one website and then try them out at all the others. Compromised accounts are sold on to people who exploit them in various ways. Primary email accounts often have recovery information for other accounts, including bank accounts if the attacker is lucky. They can also be used for scams such as the stranded traveler, where the victim emails all their friends saying they've been robbed in some foreign city and asking for urgent financial help to pay the hotel bill. If all else fails, they can be used to send spam.

2.3.1.4 Spam senders

Spamming arrived on a small scale when the Internet opened to the public in the mid-1990s, and by 2000 we saw the Earthlink spammer making millions from sending phishing lures. By 2010 spam was costing the world's ISPs and tech companies about \$1bn a year in countermeasures, but it earned its operators perhaps one percent of that. The main beneficiaries may have been webmail services such as yahoo, hotmail and gmail, which can operate better spam filters because they see so much of the world's email traffic; during the 2010s, hundreds of millions of people switched to using their services.

Spam is now a highly specialised business, as getting past modern spam filters requires a whole toolbox of constantly-changing tricks. If you want to use spam to install ransomware, you're better off paying an existing service than trying to learn it all from scratch. Some spam involves industrial-scale email compromise, which can be expensive; some \$350m was knocked off the \$4.8bn price at which Yahoo was sold to Verizon after a bulk compromise [560]. Some

spammers also operate pay-per-install services, where they might charge \$10-15 per thousand machines infected in the USA and Europe, and perhaps \$3 for Asia. A lot of pay-per-install is done through phishing campaigns, though some uses malvertising on porn sites where users are told to install a ‘new video codec’ to access a movie.

2.3.1.5 Cashout operators

Back in the twentieth century, people who stole credit card numbers would have to go to the trouble of shopping for goods and then selling them to get money out. Nowadays there are specialists who will buy compromised bank credentials on underground markets and do this. The prices show where the real value lies in the criminal chain; a combination of credit card number and expiry date sells for under a dollar, and to get into the single dollars you need a CVV, the cardholder’s name and address, and more.

Cashout techniques change every years or so, as paths are discovered through the banking system’s money-laundering controls, or ways to exploit firms that are outside the money-laundering system. Some cashout firms organise *mules* to whom they transfer some of the risk. Back in the mid-2000s, mules could be drug users who would go to stores and buy goods with stolen credit cards; then there was a period when unwitting mules were recruited by ads promising large earnings to ‘agents’ who would represent foreign companies and remit stolen funds through their personal bank accounts. Then the laundrymen used Russian banks in Latvia, to which Russian mules would turn up to get cash withdrawals. Then an unlicensed digital currency service based in Costa Rica, Liberty Reserve, was all the rage until it was closed down and its founder arrested in 2013. Bitcoin the took over but then its popularity with the cybercrime community tailed off as its price became more volatile, and as the US Department of the Treasury started arm-twisting bitcoin exchanges into identifying their customers. Since 2015–16, a growing number of ransomware variants demand Amazon gift vouchers.

As with spam, cashout is a constantly evolving attack-defence game. We monitor it and analyse the trends using CrimeBB, a database we’ve assembled of over 50 million posts in underground hacker forums where cybercriminals buy and sell services including cashout [1079].

2.3.2 Attacks on banking and payment systems

Attacks on card payment systems started with lost and stolen cards, with forgery at scale arriving in the 1980s; the dotcom boom ramped things up further in the 1990s as many businesses started selling online with little idea of how to detect fraud; and it was card fraud that spawned underground markets in the mid-2000s as criminals sought ways to buy and sell stolen card numbers as well as related equipment and services.

Another significant component is pre-issue fraud, known in the USA as ‘*identity theft*’ [481], where criminals obtain credit cards, loans and other assets in your name and leave you to sort out the mess. I write ‘identity theft’ in quotes

as it's actually the old-fashioned offence of impersonation. Back in the twentieth century, if someone went to a bank, pretended to be me, borrowed money from them and vanished, then that was the bank's problem, not mine. In the early twenty-first, banks took to claiming that it's your identity that's been stolen rather than their money [1244]. There is less of that liability dumping now but the FBI still records much cybercrime as 'identify theft' which helps keep it out of the mainstream US crime statistics.

The card fraud ecosystem is now fairly stable. Surveys in 2011 and 2019 show that while card fraud doubled over the decade, the loss fell slightly as a percentage of transaction value [66, 67]; the system has been getting more efficient as it grows. Many card numbers are harvested in hacking attacks on retailers, which can be very expensive for them once they've paid to notify affected customers and reimburse banks for reissued cards. As with the criminal infrastructure, the total costs may be easily two orders of magnitude greater than anything the criminals actually get away with.

Attacks on online banking ramped up in 2005 with the arrival of large-scale phishing attacks; emails that seemed to come from banks drove customers to imitation bank websites that stole their passwords. The banks responded with techniques such two-factor authentication, or the low-cost substitute of asking for only a few letters of the password at a time; the crooks' response, from about 2009, has been credential-stealing malware. Zeus and later Trojans lurk on a PC until the user logs on to a bank whose website they recognise; they then make payments to mule accounts and hide their activity from the user – the so-called 'man-in-the-browser attack'. (Some Trojans even connect in real time to a human operator.) A number of criminal gangs have been broken up, and people arrested, for operating these scams, but they continue to net in the low billions a year worldwide.

Firms also have to pay attention to business email compromise, where a crook compromises a business email account and tells a customer that their bank account number has changed; or where he impersonates the CEO and orders a financial controller to make a payment; and social engineering attacks by people pretending to be from your bank who talk an employee into releasing a code to authorise a payment. Most targeted attacks on company payment systems can in theory be prevented by the control procedures that most large firms already have, and so the typical target is a badly-run large firm, or a medium-sized firm with enough money to be worth stealing but not enough control to lock everything down.

I'll discuss the technicalities of such frauds in the chapter on Banking and Bookkeeping, along with a growing number of crimes that directly affect only banks, their regulators and their retail customers. I'll also discuss cryptocurrencies, which facilitate cybercrimes from ransomware to stock frauds.

2.3.3 Internal fraud

Short forward reference to Banking and Bookkeeping?

2.3.4 Sectoral cybercrime ecosystems

A number of sectors other than banking have their own established cybercrime scenes. One example is travel fraud. There's a whole ecosystem of people who sell fraudulently obtained air tickets, which are sometimes simply bought with stolen credit card numbers, sometimes obtained directly by manipulating or hacking the systems of travel agents or airlines, sometimes booked by corrupt staff at these firms, and sometimes scammed from the public directly by stealing their air miles. The resulting cut-price tickets are sold directly using spam or through various affiliate marketing scams. Some of the passengers who use them to fly know they're dubious, while others are dupes – which makes it hard to deal with the problem just by arresting people at the boarding gate. (The scammers also supply tickets at the last minute, so that the alarms are usually too late.) For an account and analysis of travel fraud, see Hutchings [674]. An increasing number of other business sectors are acquiring their own dark side, and I will touch on some of them in later chapters.

2.3.5 CEO crimes

Companies attack each other, and their customers too. From the 1990s, printer vendors have used cryptography to lock their customers in to using proprietary ink cartridges, as I describe in section 22.6, while companies selling refills have been breaking the crypto. Games console makers have been playing exactly the same game with aftermarket vendors. The US courts decided in the Lexmark case that this was fine: the incumbent could hire the best cryptographers they could find to lock their products, while the challenger could hire the best cryptanalysts they could find to unlock them. Here, the conflict is legal and open. As with state actors, corporates sometimes assemble teams with multiple PhDs, millions of dollars in funding, and capital assets such as electron microscopes. Both our hardware lab and our NGO activities have on occasion received funding from such actors.

Not all corporate attacks are conducted as openly. Perhaps the best-known covert hack was by Volkswagen on the EU and US emissions testing schemes; diesel engines sold in cars were programmed to run cleanly if they detected the standard emission test conditions, and efficiently otherwise. For this, the CEO of VW was fired and indicted in the USA (to which Germany won't extradite him), while the CEO of Audi was fired and jailed in Germany [788]. VW has set aside €25bn to cover criminal and civil fines and compensation.

Sometimes products are designed to break whole classes of protection system, an example being the overlay SIM cards described later in the chapter on Banking and Bookkeeping. These are SIM cards with two sides and only 160 microns thick, which you stick on top of the SIM card in your phone to provide a second root of trust; they were designed to enable people in China to defeat the high roaming charges of the early 2010s. The overlay SIM essentially does a man-in-the-middle attack on the real SIM, and can be programmed in Javacard. A side-effect is that such SIMs make it really easy to do some types of bank fraud.

So when putting together the threat model for your system, stop and think

what capable motivated opponents you might have among your competitors, or among firms competing with suppliers on which products you depend. The obvious attacks include industrial espionage, but nowadays it's much more complex than that.

2.3.6 Whistleblowers

Intelligence agencies, and secretive firms, can get obsessive about 'the insider threat'. But let's turn it round and look at it from the other perspective – that of the whistleblower. Many are trying to do the right thing, often at a fairly mundane level such as reporting a manager who's getting bribes from suppliers or who is sexually harassing staff. Even then, they often lose because of the power imbalance; they get fired and the problem goes on. Many security engineers think the right countermeasure to leakers is technical, such as data loss prevention systems, but robust mechanisms for staff to report wrongdoing – including ethical concerns about management decisions – are usually more important. They are also usually done badly. Internal whistleblowing mechanisms are often an afterthought and provide a complaint path that leads to HR rather than to the board's audit committee; and external mechanisms may be little better. One big service firm ran a "Whistle-blowing hotline" for its clients in 2019; but the web page code has trackers from LinkedIn, Facebook and Google, who could thus identify unhappy staff members, and also javascript from CDNs, littered with cookies and referrers from yet more IT companies. At the top end of the ecosystem, some newspapers offer ways for whistleblowers to make contact using encrypted email. But the mechanisms tend to be clunky and the web pages that promote them do not really educate potential leakers about either the surveillance risks, or the operational security measures that might mitigate them.

This is at heart a policy problem rather than a technical one. Even in the case of Ed Snowden, there should have been a robust way for him to report unlawful conduct by the NSA to the appropriate arm of government, perhaps a Congressional committee. But he knew that a previous whistleblower, Bill Binney, had been arrested and harassed after trying to do that. In hindsight, that aggressive approach was unwise, as President Obama's NSA review group eventually conceded. At the less exalted level of a commercial firm, if one of your staff is stealing your money, and another wants to tell you about it, you'd better make that work.

2.4 The Swamp

Our third category is abuse, by which we usually mean offences against the person rather than against property. These range from cyber-bullying at schools all the way to state-sponsored Facebook advertising campaigns that get people to swamp legislators with death threats. I'll deal first with offences that scale, including political harassment and child sex abuse material, and then with offences that don't ranging from school bullying to intimate partner abuse.

2.4.1 Hacktivism and hate campaigns

Propaganda and protest evolved as technology did. Ancient societies had to make do with epic poetry; cities enabled people to communicate with hundreds of others directly, by making speeches in the forum; the invention of writing enabled a further scale-up. The spread of printing in the sixteenth century led to wars of religion in the seventeenth, daily newspapers in the eighteenth and mass-market newspapers in the nineteenth. Activists learned to compete for attention in the mass media, and honed their skills as radio and then TV came along.

Activism in the Internet age gets much of its traction from online hate campaigns and radicalisation. Strikes don't seem to work so well; if you delete your Facebook account in protest, people mostly don't notice. But if you can motivate hundreds of people to send angry emails or tweets, that can get a result. A company, or individual, on the receiving end of such a campaign can have a real problem. Denial-of-service attacks can interrupt operations and doxing can do real brand damage as well as causing distress to executives and staff.

Activists vary in their goals, in organisational coherence and in the extent to which they'll break the law. There's a whole spectrum, from whistleblowers who go to respectable newspapers, through enthusiasts who harass people behind the mild anonymity of Twitter accounts, to more hard-core types who end up in jail for terrorist offences. The Climategate scandal, described in 2.2.5 above, may be an example of doxing by a hacktivist. My own earliest experience with hacktivism was with the Animal Liberation Front which targeted my university in 2003 because of plans to build a monkey house, for primates to be used in research. The online component consisted of thousands of emails sent to staff members with distressing images of monkeys with wires in their brains; this was an early example of 'brigading', where hundreds of people gang up on one target online. Back then, we dealt with the online attack easily enough by getting their email accounts closed down. But they persisted with physical demonstrations and media harassment; our Vice-Chancellor decided to cut her losses, and the monkey house went to Oxford instead. Some of the leading individuals were later jailed for terrorism offences following assaults on the staff of a local pharmaceutical testing company and after bombs were placed under the cars of medical researchers [15].

Online shaming has become popular as a means of protest. It can be quite spontaneous, with a flash mob of vigilantes forming when an incident goes viral. An early example happened in 2005 when a young lady in Seoul failed to clean up after her dog defecated in a subway carriage. Another passenger photographed the incident and put it online; within days the 'dog poo girl' had been hounded into hiding, abandoning her university course [314]. There have been many other examples.

We discovered the power of platforms such as Twitter in Gamergate, a storm sparked by abusive comments about a female game developer made publicly by a former boyfriend in August 2014, and cascading into a torrent of misogynistic criticism of women in the gaming industry and of feminists who had criticised the industry's male-dominated culture. A number of people were doxed, SWATted, or hounded from their homes [1364]. The harassment was coordinated on

anonymous message boards such as 4Chan and the attackers would gang up on a particular target – who then also got criticised by mainstream conservative journalists [829]. The movement appeared leaderless and evolved constantly, with one continuing theme being a rant against ‘social justice warriors’. It may even have contributed to the development of the alt-right movement which influenced the 2016 election two years later.

A growing appreciation of the power of angry mobs is leading politicians to try to stir them up, at all levels from people trying to do down their rivals for local political office to nation states advertising covertly to swing rival states’ elections. Companies are targeted less frequently but it does happen. Meanwhile the social-media companies are under pressure to censor online content, and as it’s hard for an AI program to tell the difference between a joke, abuse, a conspiracy theory and information warfare by a foreign government, they end up having to hire more and more moderators. I will return to the law and policy aspects of this in 24.4 below.

2.4.2 Child sex abuse material

When the Internet came to governments’ attention in the 1990s and they wondered how to get a handle on it, the first thing to be regulated was images of child sex abuse (CSA), in the Budapest Convention on 2004. We have little data on the real prevalence of CSA material as the legal restrictions make it hard for anyone outside of law enforcement to do any research. In many countries, the approach to CSA material has less focus on actual harm reduction than it deserves. Indeed, many laws around online sexual offences are badly designed, and seem to be driven more by exploiting outrage than by minimising the number of victims and the harm they suffer. CSA gives us a case study on how not to do online regulation because of forensic failures, takedown failures, weaponisation and the law-norm gap.

The most notorious forensic failure was Britain’s Operation Ore, which I describe in more detail in 24.5.3; briefly, several thousand men were arrested on suspicion of CSA offences after their credit card numbers were found on an abuse website, and perhaps half of them turned out to be victims of credit card fraud. Hundreds of innocent men had their lives ruined. Yet nothing was done for the child victims in Brazil and Indonesia, and the authorities are still nowhere near efficient at taking down websites that host CSA material. In most countries, CSA takedown is a monopoly of either the police, or a regulated body that operates under public-sector rules (NCMEC in the USA and the IWF in the UK), and takes from days to weeks; things would go much more quickly if governments were to use the private-sector contractors that banks use to deal with phishing sites [675]. The public-sector monopoly stems from laws that make the possession of child sex abuse material a strict-liability offence, which not only makes it hard to deal with using the usual abuse mechanisms. It also allows it to be weaponised: protesters can send it to targets and then report them to the police. It also makes it difficult for parents and teachers to deal sensibly with incidents that arise with teens using dating apps or having remote relationships. The whole thing is a mess, caused by legislators wanting to talk tough without understanding the technology.

The latest problem is an emerging law-norm gap with the growth in popularity of sexting among teenagers. Like it or not, sending intimate photographs to partners (real and intended) has become normal behaviour for teens in many countries, despite the fact that possessing an intimate photo of anyone under 18 can now result in a prison sentence. Teens laugh at lectures from schoolteachers to not take or share such photos, but the end result is real harm. Kids may be tricked or pressured into sharing photos of themselves, and even if the initial sharing is consensual, the recipient can later use it for blackmail or just pass it round for a laugh. This leads to general issues of bullying and more specific issues of intimate partner abuse.

2.4.3 School and workplace bullying

Online harassment and bullying are a fact of life in modern societies, not just in schools but in workplaces too, as people jostle for rank, mates and resources. From the media stories of teens who kill themselves following online abuse, you might think that cyber-bullying is now most of the problem bullying, but the figures show that it's less than half. An annual UK survey discloses that about a quarter of children and young people are constantly bullied (13% verbal, 5% cyber and 3% physical) while about half are bullied sometimes (24%, 8% and 9% respectively) [415]. The only national survey of all ages of which I'm aware is the French national victimisation survey, which since 2007 has collected data not just on physical crimes such as burglary and online crimes such as fraud, but on harassment too [1051]. This is based on face-to-face interviews with 16,000 households and the 2017 survey reported two million cases of threatening behaviour, 7% were made on social networks and a further 9% by phone.

There is also talk in the media of a rise in teen suicide which some commentators link to social media use. Thankfully, the OECD mortality statistics show that this is also untrue: suicides among 15–19 year olds have declined slightly from about 8 to about 7 cases per 100,000 over the period 1990–2015 [1061].

2.4.4 Intimate partner abuse

Just as I ended the last section by discussing whistleblowers – the insider threat to companies – I'll end this section with intimate partner abuse, the insider threat to families and individuals. Gamergate may have been a flashbulb example, but protection from former intimate partners is a real problem that exists at scale – with about half of all marriages ending in divorce, and not all breakups being amicable. Intimate partner abuse has been suffered by 27% of women and 11% of men. Stalking is not of course limited to former partners. Celebrities in particular can be stalked by people they've never met – with occasional tragic outcomes, as in the case of John Lennon. But former partners account for most of it, and law enforcement in most countries have historically been reluctant to do anything effective about them.

One subproblem is the publication of non-consensual intimate photographs, once called 'revenge porn' – until California Attorney-General Kamala Harris argued that this is cyber-exploitation and a crime. Her message got through to

the big service firms who since 2015 have been taking down such material on demand from the victims [1218]. This followed an earlier report in 2012 where Harris documented the increasing use of smartphones, online marketplaces and social media in forcing vulnerable people into unregulated work including prostitution – raising broader questions about how technology can be used to connect with, and assist, crime victims [623].

The problems faced by a woman leaving an abusive and controlling husband are among the hardest in the universe of information security. All the usual advice is the wrong way round: your opponent knows not just your passwords but has such deep contextual knowledge that he can answer all your password recovery questions. There are typically three phases: a physical control phase where the abuser has access to your device and may install malware, or even destroy devices; a high-risk escape phase as you try to find a new home, a job and so on; and a life-apart phase when you might want to shield location, email address and phone numbers to escape harassment, and may have lifelong concerns. It takes seven escape attempts on average to get to life apart, and disconnecting from online services can cause other abuse to escalate. After escape, you may have to restrict childrens' online activities and sever mutual relationships; letting your child post anything can leak the school location and lead to the abuser turning up. You may have to change career as it can be impossible to work as a self-employed professional if you can no longer advertise.

To support such users, responsible designers should think hard about usability during times of high stress and high risk; they should allow users to have multiple accounts; they should design things so that someone reviewing your history should not be able to tell you deleted anything; they should push 2-factor authentication, unusual activity notifications, and incognito mode. They should also think about how a survivor can capture evidence for use in divorce and custody cases and possibly in criminal prosecution, while minimising the trauma [906]. But that's not what we find in real life. Many banks don't seem to want to know about disputes or financial exploitation within families. Most spyware vendors are eager to sell to controlling husbands. And then there's the Absher app, which enables men in Saudi Arabia to control their women; its availability in app stores has been controversial elsewhere.

2.5 Summary

The systems you build or operate can be attacked by a wide range of opponents. It's important to work out who might attack you and how, and it's also important to be able to figure out how you were attacked and by whom. Your systems can also be used to attack others, and if you don't think about this in advance you may find yourself in serious legal or political trouble.

In this chapter I've grouped attackers under three themes: the spooks, the crooks and the swamp. Intelligence and law enforcement agencies typically use the analysis of traffic data when hunting, and targeted collection for gathering, whether via malware, deception or coercion. Both spooks and crooks use malware to establish botnets as infrastructure. Crooks typically use opportunistic collection for mass attacks and deception for targeted work; as with the spooks,

spear-phishing is the weapon of choice. There are also cybercrime ecosystems attached to specific business sectors; basically, crime will evolve where it can scale. As for the swamp, the weapon of choice is the angry mob, wielded nowadays by states, activist groups and even individual orators. There are many ways in which abuse can scale, and when designing a system you need to work out how crimes against it, or abuse using it, might scale. It's not enough to think about usability; you need to think about abusability too.

If you're defending a company of any size, you'll see enough machines on your network getting infected, and you need to know whether they're just zombies on a botnet or part of a targeted attack. So it's not enough to rely on patching and antivirus. You need to watch your network and keep good enough logs that when an infected machine is spotted you can tell whether it's a kid building a botnet or a targeted attacker who responds to loss of a viewpoint with a scramble to develop another one. You need to make plans to respond to incidents, so you know who to call for forensics – and so your CEO isn't left gasping like a landed fish in front of the TV cameras. You need to think systematically about your essential controls: backup to recover from ransomware, payment procedures to block business email compromise, and so on. If you're advising a large company they should have much of this already, and if a small company you need to help them figure out how to do enough of it.

The rest of this book will fill in the details.

2.6 Research Problems

Until recently, research on cybercrime wasn't really scientific. Someone would get some data – often under NDA from an anti-virus company – work out some statistics, write up their thesis, and then go get a job. The data were never available to anyone else who wanted to check their results or try a new type of analysis. Since 2015 we've been trying to fix that by setting up the Cambridge Cybercrime Centre, where we collect masses of data on spam, phishing, botnets and malware as a shared resource for researchers. We're delighted for other academics to use it. If you want to do research on cybercrime, call us.

We also need something similar for espionage and cyber warfare. People trying to implant malware into control systems and other operational technology are quite likely to be either state actors, or cyber arms vendors who sell to states. The criticisms made by President Eisenhower of the 'military-industrial complex' apply here in spades. Yet not one of the legacy think-tanks seems interested in tracking what's going on. As a result, nations are more likely to make strategic miscalculations, would could lead not just to cyber conflict but the real kinetic variety too.

As for research into cyber abuse, there is now some research, but the technologists, the psychologists, the criminologists and the political scientists aren't talking to each other enough. There are many issues, from the welfare and rights of children and young people to our ability to hold fair and free elections. We need to engage more technologists with public-policy issues and educate more policy people about the realities of technology.

2.7 Further Reading

There's an enormous literature on the topics discussed in this chapter but it's highly fragmented. A starting point for the Snowden revelations might be Glen Greenwald's book *'No Place to Hide'* [592]; for surveys of cybercrime, see our 2012 paper "Measuring the Cost of Cybercrime" [66] and our 2019 follow-up "Measuring the Changing Cost of Cybercrime" [67]; and for a great introduction to the history of propaganda, see Tim Wu's *'The Attention Merchants'* [1455].