A "sophisticated attack"? Innovation, technical sophistication, and creativity in the cybercrime ecosystem

Ben Collier, University of Edinburgh and Richard Clayton, University of Cambridge

Abstract

We observe that almost every cybercrime is reported to be a "sophisticated attack" and explain how incentives align to misrepresent very run-of-the-mill events in this manner. We describe the cybercrime ecosystem, analysing the distinct parts and discussing what forms of sophistication and incentives can be found in each kind of work. We move on discuss how framing cybercrime as technically sophisticated attacks performed by skilled criminals has distorted criminological analysis and contributed to misaligned incentives within criminal justice and security policy. We conclude that the criminal justice system is aiming the wrong types of interventions at the wrong kinds of actor.

sophisticated, *adj*: ... **c:** Of equipment, techniques, theories, etc.: employing advanced or refined methods or concepts; highly developed or complicated. ... OED [81]

sophisticated, *adj*: **1:** deprived of native or original simplicity: such as **a:** highly complicated or developed. ... Merriam-Webster [64]

1 Introduction

In April 2020 a 17-year old schoolboy, operating from his bedroom in Lincolnshire in the UK, set up a fake website which impersonated Love2shop a UK website that sells gift vouchers. To attract people to his website he purchased Google adverts so that his site appeared at the top of search result pages, above the genuine site. He took down the website after a week having defrauded Love2shop of £6 000, but consumer complaints meant that he became the subject of a police investigation.

When he was arrested in August 2020 it was found that his computer had details of 12 000 credit card numbers (presumably from other scams besides Love2shop) and he owned 197 PayPal accounts that had received £323 000 between them. From this, £200 000 had been invested in cryptocurrency and by the time of his trial in October 2021, his holdings had soared to a value of over two million pounds. The schoolboy was sentenced, in line with UK practice for someone of his age and lack of previous convictions, to 150 hours of unpaid work and given a 12 month youth supervision order. The court also found that he had benefited from his crimes to the amount of £2 141 720 and relevant assets were seized [82].

This is, in our view, a pretty run-of-the-mill cybercrime, distinguished only by the fact it was successfully investigated and the culprit brought before a court. Impersonating websites is relatively trivial, using Google adverts to drive traffic is conventional business practice, opening large numbers of PayPal accounts to avoid undue notice a standard criminal tactic and failing to keep sensitive business records encrypted sadly all too common.

Clearly the police investigation was more complex than expected since the original fraud led them to discover other offences, but this does not seem especially unusual. However, the police told the press, "this was a sophisticated cyber fraud". The judge took the same view saying "You have a long-standing interest in computers. Unfortunately, you used your skills to commit a sophisticated fraud".

Workshop on the Economics of Information Security (WEIS), June 2022.

This inflation of nomenclature is widespread. In October 2015 TalkTalk, a UK ISP, was hacked and customer details were stolen. In their FAQ, issued immediately after the attack was discovered, TalkTalk said:

> We believed our systems were as secure as they could be. We work with world leading security experts and update our systems constantly. [...] Unfortunately these criminals are very smart and their attacks are becoming ever more sophisticated [88].

It subsequently emerged that the breach involved a 'SQL injection' attack that had been performed by various UK teenagers – viz the attack method was older than they were. TalkTalk was eventually fined a record-breaking £400 000 by the Information Commissioner for a "failure to implement the most basic cyber security measures [which] allowed hackers to penetrate TalkTalk's systems with ease." [36].

As a further example of exaggeration and inaccuracy, just prior to the 2019 election, a UK Labour Party website was subjected to a trivial Denial of Service attack of the type that can be purchased for a few pounds from any number of services operating openly on the clearnet. Labour's website was hosted on a specialist Denial of Service protection system, but was using the lowest (and cheapest) level of protection, and it suffered several outages [87]. Nevertheless Niall Sookoo, the party's executive director of elections and campaigns told campaigners: "Yesterday afternoon our security systems identified that, in a very short period of time, there were large-scale and sophisticated attacks on Labour Party platforms which had the intention of taking our systems entirely offline. Every single one of these attempts failed due to our robust security systems and the integrity of all our platforms and data was maintained" [15].

To be clear at the outset, we are not saying that there are no sophisticated cybercrimes, for example the Solarwinds attack involved forging SAML certificates which allowed the APT group "UNC2452/Dark Halo" access to many thousands or corporate and government websites and this was undetected for a considerable time [77]. Additionally, although not everyone would regard them as crimes, the Stuxnet malware had considerable technical sophistication (which is why it was studied so intently long before its actual purpose of disrupting Iran's nuclear programme was understood) [54]; and no-one could use any other word than sophisticated to read how the NSO Group's FORCEDENTRY attack proceeds by exploiting a PDF buffer overflow to utilise some obscure JBIG2 functionality to create a Turing complete computer out of virtual logic gates [16].

In this paper we explore in Section 2 why simple cybercrimes are being described as "sophisticated" and then in Section 3 how this is distorting our approach to dealing with them. In Section 4 we map out the different groups of cybercriminals and set out how sophisticated their activity might be. We discuss what this means for criminological theory in Section 5 and for cybercrime policy in Section 6. We then survey how others have approached these issues in Section 7 before concluding in Section 8.

2 Defender incentives for mischaracterising sophistication

It might be thought that cybercriminals always needed to employ sophisticated techniques in order to succeed. Sadly this is not the case. In 2011 the UK Government said that "80% or more of currently successful attacks exploit weakness that can be avoided by following simple best practice" [91], (a statement often misunderstood to mean that following best practice would reduce cybercrime by 80%). Despite efforts to improve cybersecurity, there is little evidence so far of any overall change [14].

Unsophisticated attacks work because almost every system is insecure at some level. The lower levels of architecture, standards, and protocols on which computer systems rely are very hard to change and inherit a range of vulnerabilities from early in their history. Fixing these often means altering basic aspects of how computers work – e.g. the 'capabilities' added by CHERI [97] – but this requires huge resources and decades of research development, implementation, and then adoption. Above the hardware and operating systems are multiple complex levels of software within which the inter-level interactions are effectively infinite. Any implementation flaw is a potential insecurity – and once the nature of that insecurity is understood, attacks can be automated and placed into the hands of pretty much anyone. As Cisco's

Chairperson put it in 2015, "There are two types of companies: those who have been hacked, and those who don't yet know they have been hacked" [19].

Almost everyone involved in cybercrime has an incentive to overstate either the impact of incidents, or their complexity, or both. The attackers want to project (to their peers, but not of course in an attributable way) an image of competence and success. Corporate victims will wish to claim that their attackers were so clever no-one could have withstood their attack and hence that they cannot be blamed for inadequate defences – hoping thereby to create the impression that civil suits for damages, or regulatory action, would be unreasonable. Even in private, the claim of sophistication may be maintained, as the company's IT Department claims that future problems are inevitable unless they receive substantial amounts of money to counter them.

If victims wish Law Enforcement to act then their losses may need to pass a predetermined threshold, which may require a creative approach to summing up costs, but it is also helpful to set out a narrative that the attack was so sophisticated that it should be investigated anyway. Law Enforcement are then incentivised to talk up the competence of the criminals to explain their lack of success in catching them, or occasionally, their own outstanding skill in making an arrest.

The wider security industry have clear incentives to talk up risk and the need for (their proprietary) sophisticated products to counter the sophisticated attacks – and individuals within the industry may find their career depends upon a regular output of blogs and BlackHat talks which detail exotic attacks; because explaining run-of-the-mill issues is too boring to make good PR for their company.

However, we are not necessarily suggesting that the commentary about the sophistication of cybercrimes is invariably a set of cynical misstatements. As Clarke's Third Law puts it: "any sufficiently advanced technology is indistinguishable from magic" [21] and many of the people who are making the claims about sophistication have very limited exposure to the reality of modern cybercrime, failing to appreciate how commonplace most of the techniques and mechanisms being employed actually are. Perhaps the most egregious example of this gap between reality and perception in recent years concerned the Governor of Missouri, who in October 2021 ordered an investigation into a reporter on charges of 'hacking' for simply viewing the source code of a (poorly secured) government web page [51].[1]

The reality is that what counts as sophisticated changes over time, as younger generations ('digital natives') use computers, smartphones and the Internet as an integral part of their everyday lives. UX (User eXperience) developments (both in licit and illicit technologies) and the circulation of 'How-to' guides have brought complex activities to the attention of many and within the competence of all. Activities that would have required substantial technical nous and persistence ten or twenty years ago, such as invoking a botnet, delivering search-based adverts for illicit services, or laundering money through cryptocurrency, are now routinely performed by neophytes.

This failure to understand the commonplace is compounded by an unnecessary emphasis by the experts on nit-picking jargon – many an unnecessary dispute over the difference between worms, trojans and viruses was resolved by calling them all malware, an etymological success that has seldom been repeated. There remains a delight in perpetuating jargon from underground Internet culture, often from more than 30 years ago (133t, p0wn, phishing), and (as we have already mentioned above) a clear tendency to get carried away by a wish to make every minor change to a crime type seem new and exciting – with its own new set of vocabulary: whaling, pharming, vishing, smishing, quishing etc.

Far from being a minor annoyance, this serves to distort understandings of cybercrime as a phenomenon, and to drive research and policy in less productive directions. The cybersecurity industry has clear incentives to inflate the apparent severity, sophistication, and 'fast-paced' nature of cybercrime – along with these narratives, minor variations in techniques are then picked up by a media well-aware of the appeal

[1] It is worth noting, however, that the US Department of Justice in May 2022 announced a new policy that "good-faith security research should not be charged" under the Computer Fraud and Abuse Act [92].

to readers of a 'novel' cybercrime development. This then drives a vicious cycle between practitioners, policymakers, law enforcement, and academia, in which severity and sophistication are invoked but rarely systematically and empirically examined outside the incentive systems of law enforcement and private security. This has important consequences for how law enforcement and the broader criminal justice system respond to cybercrime, which we discuss in the following section.

In fact, following decades of innovation, transformation, and growth, cybercrime has now largely stabilised, making up around half of all acquisitive crime. However, the underlying vulnerabilities, technologies, and infrastructures have changed little in the last decade [5]. The main novelty in recent years has been ransomware actors changing focus from victimising individuals to attacking companies, and exfiltrating sensitive data to counter backup strategies. Additionally, mass leaks of password data has lead to a rise in 'credential stuffing' (trying a username/password against a wide range of websites) rather than just guessing passwords from a list of common choices. These represent significant changes in the organisation of existing forms of crime, rather than technically novel or sophisticated developments. Most of the perceived innovations in cybercrime of recent years (and even the much-vaunted changes to cybercrime over the course of the pandemic) involve changes in volume between different crime types, minor changes to delivery methods, or alterations in business models, illicit marketing and bait strategies, with the core crime itself following largely the same pattern and the same socio-technical underpinnings and vulnerabilities as before.

Notwithstanding that some commentators may just be ill-informed, the main reason for the overwhelming majority of cybercrimes to be described as 'sophisticated' is very clearly the incentives for doing so. For over a decade misaligned incentives have been seen as a key factor in failures to defend against cybercrime [66]. In this case, however, all the incentives are aligned and although policy makers clearly understand that there is a still a need to get the basics right (to prevent 80% of cybercrime being trivial to perform) this message has to fight against the conception that nothing in this space could ever be seen as trivial.

3 Why does sophistication matter?

The level of technical sophistication of a crime is *important* and this has real relevance for criminal justice policy and practice. In this section, we briefly discuss the relevance of cybercrime sophistication in judicial and policing contexts.

3.1 Sophistication and the courts

When cybercrimes go to court, perceived sophistication can be an aggravating factor. In the UK cybercrime falls under the Computer Misuse Act – and, while technical sophistication is not directly listed as an 'aggravating factor' to be considered in sentencing, judges do take this into account (as we noted above). It also forms one of the aspects in determining the Public Interest component used by the Crown Prosecution Service when deciding whether to prosecute:

> "The level of sophistication used, particularly sophistication used to conceal or disguise identity (including masquerading as another identity to divert suspicion)" [27]

The approach is to assume that sophistication generally corresponds to a higher perceived seriousness and culpability of the offence, as more technical offences are assumed to have higher rates of foreknowledge, planning, organisation, and commitment on the part of the offender, and to be more generally "involved, with a higher degree of mens rea and likelihood of re-offense".

In a US legal context, sophistication is more explicitly brought in at the sentencing process for cybercrimes, acting as an *enhancement* (i.e. a factor which increases the mandatory minimum sentence applicable). As Kessler sets out in a recent overview paper on US cybercrime sentencing [46], for Section 1030 (cybercrime) offences, enhancements from Section 2B1.1(b)(10) will apply in many cases, including:

> [use of] "sophisticated means" – and "especially complex or especially intricate offence conduct pertaining to the execution or concealment of an offence"

And similarly, enhancements from Section 3B1.3 apply to many Section 1030 offences:

> "the use of a special skill... not possessed by members of the general public and usually requiring substantial education, training, or licensing"

This is drawn from US legislation regarding fraud prosecutions, in which 'sophisticated means' was brought in during 1998 as a way of distinguishing particularly harmful and difficult-to-detect offences. In practice, the application at sentencing of the 'sophisticated means' enhancement makes imprisonment of some length likely to occur for very low-level fraud, increases sentences by around 6 months for minor fraud, and prolongs sentence length by up to 20 or 30 months for major fraud [12]. The rationale behind this enhancement operates on an economistic deterrence logic: more sophistication in the commission of offences is generally deemed to correspond to a lower likelihood of detection and arrest. Thus, harder-to-catch crimes accrue increased punitive sanctions in order to attempt to increase the perceived risk in the risk-benefit calculus of the prospective offender. A separate justification operates on a moral logic – that sophisticated crimes are likely to involve more premeditation and hence culpability [12, 17].

Interestingly, the percentage of frauds assessed by the courts as sophisticated more than tripled from 2.9% in 2005 to 11.7% in 2013 [12]. Baer's scholarship on the enhancement argues that, rather than fraud crimes themselves becoming generally more sophisticated, a process of net-widening or 'sentencing creep' has occurred. The bar for sophistication has progressively dropped over time, with advancements in technology significantly simplifying means which previously would have required substantial sophistication to execute [12, 17]. Accordingly, the judicial practices of a range of actors have progressively moved to encourage the use of this enhancement more and more broadly over a period of several years [12]. In our view, a very similar process appears to be occurring for cybercrime.

Although legal scholarship has for some time recognised that there is a distinction between tool-creating and tool-using forms of cybercrime in terms of sophistication [45], sophistication is itself a slippery concept, and, as we discuss in this paper, the rise of volume cybercrime (and an increasingly stabilised cybercrime ecosystem) has further complicated assessments of sophistication by non-technical experts.

3.2 Sophistication and the police

For 'volume' cybercrime, the police (in the UK and elsewhere) not only arrest and prosecute those found to have committed particular crimes, but can also invoke a range of preventative and pro-active forms of intervention, such as 'restorative' reparation of harm caused, or attempts to encourage a a pro-social identity and a change in behaviour.

The police use 'diversion programmes', although unlike most diversionary work, police attempts at diversion generally focus more on 'net-widening' than true diversion, which aims at keeping people out of the criminal justice system and preventing from them being drawn further into law enforcement contact [22, 63]. The emphasis within policing is on preventative rationales, generally involving some form of surveillance in order to identify those at most risk of committing particular kinds of crime, and then subsequently the delivery of targeted interventions. The perception that cybercrimes involve significant technical skill has generally meant that the profile of an 'at-risk' young person includes an interest in computer science and technical skill with computers, along with other factors, such as an autism diagnosis [79]. However, there is no evidence linking ASD conditions with a heightened risk of engaging in cybercrime, and this hypothecated relationship itself is held to operate through increased focus and hence ability to develop technical skills, which may also be a misapprehension [55, 58].

Targeted interventions are focused on young people who have already begun to commit more serious forms of crime, although they may not have received a custodial sentence. These involve reparative measures where harm has been done, along with additional support and training. The aim is to repurpose deviant skills towards positive ends (the poacher-turned-gamekeeper hypothesis), which has been a particularly

influential perspective for cybercrime preventative programmes, such as encouraging young people towards a career in cybersecurity, or mentorship programmes involving industry [79].

Decisions around enforcement and punishment still tend to be shaped by the idea that those who commit cybercrime are somehow *different* from those who commit other forms of crime – possessed of important marketable skills which could be repurposed to socially positive ends. While a diversion is no doubt a more socially positive approach when compared to punitive imprisonment-based sentences, assumptions around the skills of young people involved in cybercrime may mean that diversionary or desistance programmes are assuming substantially more technical skill than may actually be evident. Few of the people approached by the police for these diversionary programmes have the skills required by the cybersecurity industry – or necessarily any interest in learning them. Conversely, the skills developed by most people involved in cybercrime are those required to run a small business – marketing, customer management, and the development of a profitable business model [62, 67].

4 Mapping cybercrime as an ecosystem

We now discuss groups doing distinct kinds of work within the cybercrime ecosystem. We treat these groups of workers in a roughly descending order of *technical* 'sophistication', but outline the distinct forms of sophistication, innovation, and creativity which are evident in each. For each grouping, we discuss the kind of work involved, the role the group plays in the cybercrime ecosystem, what levels of technical or other sophistication might be expected, and, finally, how this shapes the incentives and other economic factors at play.

4.1 Hacking – basic research and paradigm shifts

Popular culture considers all cybercriminals to be hackers – experts in computer systems who are continuously identifying and deploying new attack methods. In reality, the number of people with the resources, organisational support, creativity, skills, and training to innovate in this manner is tiny.

The reverse, that all hackers are cybercriminals, is vigorously disputed by those who point to the intellectual appeal of a deep understanding of systems and bending them to their will – but the vocabulary battle is pretty much over and so we have to describe the non-criminal activities as "ethical hacking", or, formerly, "white-hat hacking" in order to avoid complete confusion.

Research into IT system vulnerabilities and their disclosure is a part of a widely-accepted theory, practice, and ideology of security [98]. It is hypothecated that there is an arms race between competing nation states and other actors with the capabilities to do this kind of research. However, unlike the nuclear arms race, the exploits that are discovered are actually used day to day, and once 'in the wild' they are inevitably taken up by malicious actors. Hence, the theory goes, we need to race ahead and make the vulnerabilities and mitigations public through a vulnerability disclosure infrastructure and thereby make the Internet a more secure place. However, whether this theory can be shown to be true seems to depend on the assumptions one makes and how faithful the modelling is to the real world [9, 10, 35, 75, 76, 78].

We can conceive broadly of two kinds of novelty or innovation at this level. Much as in other forms of primary industrial research, most 'discoveries' or advances are piecemeal – rather than entirely new kinds of vulnerability. The day-to-day work tends to involve minor progress, with a focus on new occurrences of bugs that follow well-established patterns, new ways (or new places) to exploit existing vulnerabilities, or working through the potential 'attack space' of a novel vulnerability discovered by another group. While some of this work still requires substantial manual effort and creative practices, there is a large automated component, involving, for example, scanning through every possible call on an API in different contexts and measuring behaviour, or developing security frameworks for particular organisational contexts [99]. Much of this work cannot really be seen as technically sophisticated although this does not stop it being skilled or valuable, in many senses, as it forms the bedrock of much industry practice – given the importance of patching or mitigating relatively well-known vulnerabilities.

The second type of attack 'discovery' can be conceived of as a 'paradigm shift' [28] – one which opens up an entirely new attack space and hence transforms the landscape. Genuinely novel attack types are very rarely created by criminals and first identified in the wild; they are usually found by academics and security professionals (who publish the details) or nation states (who do not). The discoveries are essentially 'science' and the attacks involved are almost invariably 'sophisticated'. Rather than 'lightning flash' innovation, these developments are themselves often the result of decades of piecemeal advances [83].

This again mirrors the literature on paradigm shifts in industry – where a novel discovery or innovation upends the landscape, but often requires substantial investment even after discovery before it becomes competitive. As with industrial research, it takes substantial effort and expense to plough new territory within a novel paradigm, as actors learn the ropes, implement tools and develop processes. Despite the immense difficulty a novel paradigm faces at these early stages in becoming competitive with existing technologies (often called the sailing ship effect [39]) there is also a well-developed literature on the benefits which accrue to those who switch early and are able to establish themselves [29, 11, 31]).

Once discovered, new vulnerabilities can become widely known very quickly – academic (and industry) publishing incentives, media ecosystems and the complexities of vulnerability disclosure often combine to make a new attack (with explanatory website, cutesy name and logo) front page news. Most of the truly novel attacks – e.g. rowhammer [47] – are never used at any scale in the wild or it may take years before people understand them sufficiently to turn them into something applicable to the cybercrime ecosystem [3]. However, in the, thankfully, fairly rare cases when the 'attack' consists of pointing out a trivial-to-exploit vulnerability in a widely deployed software component then it may be only a few hours before very large numbers of systems are being successfully attacked and real damage is being done – examples are 'Heartbleed' (a 2014 SSL vulnerability that exposed the contents of server memory) [70] and 'Log4jShell' (a 2021 vulnerability in a Java logging component that allowed attackers to run their code on servers) [60].

Most of the basic research into attack methods is not of itself illegal, and the people who do it rarely appear in court when the law is broken. More common is the threat of civil action by the vendors of devices whose customers suddenly find themselves at risk [32]. This form of hacking is undoubtedly sophisticated in the technical sense – requiring deep knowledge of technical systems, architecture design, and principles of computing. While creativity is undoubtedly important, as with any primary scientific research, the environment for innovation relies on substantial resources, time, patience, and supportive teams to help write research grant proposals and carry out administrative work, and purchase expensive tools. Despite the counter-cultural ethics and aesthetics which draw some people into this work and which still animate the lives of conferences like PETS, WWW, or BlackHat, this kind of hacking has more or less entirely moved out of the scrappy spaces of bedroom experimentation and online deviance. Instead, it is largely the preserve of nation state espionage agencies, academic research institutes, and some private security companies.

The incentives and economic forces driving vulnerability research and disclosure are complex. However, we argue, the vulnerability economy is only one part of a wider ecosystem, the other components of which are driven by their own economic, technical, and cultural forces. As we discuss now, these vulnerabilities and mitigations are not ready-to-use at the moment of disclosure, instead providing the raw materials for communities of security practice (both offensive and defensive) and for the tools which underpin cybercrime and security.

4.2 Skilled actors, penetration testers and APTs

Although vulnerability research is a crucial central component of the wider cybersecurity ecosystem, it is also very much its own domain: time-consuming, specialised, and to most, fairly arcane. As a result, it is rare to see primary vulnerability researchers (even those working for nation state attack groups) using their attacks in practice themselves. While a vulnerability can be successful exploited under lab conditions, using it in practice to comprise computer systems involves rather a different kind of sophistication – the skills of the information security professional.

Many large companies attempt to keep their systems secure by 'penetration testing' – a 'red team' is permitted to use any and all techniques to break into the company's systems, whilst a 'blue team' is responsible for defending the systems, by keeping software up-to-date, ensuring configurations are secure and monitoring for unusual events [85]. The company may also extend the reach of their red team by operating a 'bug bounty' system where they will pay third parties who report insecure systems to them.

At the same time, the systems may be being attacked by criminals or by so-called APT groups. Those who succeed in carrying out these attacks generally have a similar skillset to 'red teamers', involving a knowledge of attack techniques and tools, practical skills in intrusion, and a wealth of tacit knowledge gained through experience. On the criminal side, threat actors generally work in teams or gangs, often associated with existing networks who facilitate cashing out, supportive activities (such as maintaining infrastructure and dealing with victims through elaborate 'customer service' operations) and monetisation. APT stands for Advanced Persistent Threat but we use the term here, as do many others, as a catchall for groups operating in the interests of states rather than specifically for their own enrichment. Although separated in terms of capabilities and explicit state support, in fact these groups are sometimes more connected than one might assume, with some criminal groups operating with tacit state support (or at least disinterest) as long as they focus their activities abroad.

All these groups are almost invariably using a playlist of known exploits in attempting to compromise their current target. They may string these exploits together in innovative ways, but in essence the task being performed is derivative of other people's innovations. Indeed, it is a commonplace to stress that the real threat posed by APT groups is hardly 'advanced' at all, but is dangerous only by its persistence – and that sooner or later a defensive lapse will give them success.

Penetration testing is often promoted (you can get paid for doing exactly what you currently do for free, and as a bonus we won't put you prison) to youngsters at an early stage of their criminal career [79]. However, there is somewhat of a difference in the skills required of an employee who is going to create a report listing all the flaws that a company needs to fix, and a cybercriminal who can enjoy considerable success by running a small number of exploits against a large number of targets, some percentage of which will be vulnerable. Hence, the youngsters may require considerable training before they are useful in their new role. In fact the defenders (the 'blue teams' mentioned above) are likely to have far more training, skill, resources than most of the attackers, but asymmetries caused by huge attack surface and the reality that attackers 'only have to be lucky once' will often make this irrelevant.

There is considerable sophistication at this level of practice, focused on professional knowledge rather than necessarily the arcane IT systems expertise which underpins exploit development. Although practitioners often have a fairly deep knowledge of computer systems, with a good understanding of the functions and interrelations between different kinds and layers of system, many of the skills involved are focused on project management, organisation, and learning practice frameworks (either in the context of a corporate environment, or a deviant one). The skills involved in carrying out or defending against an attack involve both social and technical capabilities, and this is where we do observe creativity, with forms of social engineering and technical vulnerabilities combined in novel and creative ways [85, 32].

Although this group are not developing sophisticated new types of attack, they may demonstrate considerable sophistication when they create PoCs[2] to automate their work. Even more sophistication is required to 'reverse engineer' a software patch to deduce what vulnerability is being fixed and only then generate the PoC that can be used to check for unpatched devices. Crimeware authors and APT actors naturally move beyond a mere PoC, which just reports that a device is vulnerable, into developing code that can actually exploit that vulnerability.

[2]Proof of Concept programs that demonstrate that an exploit that has been documented in general terms (perhaps in a CVE document) will work against a particular device.

This part of the ecosystem is all about the consumption rather than the creation of exploits. Economic success is closely related to the effort expended in looking for vulnerabilities, or the skill of sales and marketing in acquiring new customers who will pay for the search. A quick demonstration that a company is insecure performed by a really skilled pentester may lead to a lucrative long-term engagement where less-skilled operatives can provide a regular report that provides a little real security and much peace-of-mind.

4.3 Tool builders and infrastructure providers

A well-established sector of the cybercrime ecosystem is the provision of tools to underpin 'crime as a service' models. Some tool builders work their way down lists of vulnerabilities and package them up into specialist programs that will permit others to break into machines [49]. Other tool builders may create 'phishing kits' that can be deployed in moments to provide webforms to capture credentials [26]; they may provide turnkey websites for High Yield Investment Programs (which are Ponzi schemes) [68]; or they may operate booters that provide low cost denial of service attacks [41].

Sometimes tools shade over into infrastructure: underground forums or marketplaces may well use off-the-shelf software for basic messaging and discussion, but the system may be customised to provide reputation indicators or payment escrow functionality [1, 93]. In other cases, what is made available is clearly infrastructure: Tor hidden services are difficult for Law Enforcement to locate if the onion routing protocols are correctly implemented and care taken with the website content [71]; for 'bullet-proof hosting' the servers and network provision are standard, but the provider actively thwarts attempts to identify users or refuses to close down systems [61]. Infrastructure brings with it its own kinds of work – reliant on administration and maintenance, and emphasising stability. Rather than attempting to break systems in novel ways, the most useful 'hacks' for many of these actors are used to keep rickety illicit infrastructure running, either through automation or substantial manual maintenance [23].

Tools may have many facets to them. Recently we have seen the provision of 'ransomware as a service' to those who wish to extort victim companies whose systems have been penetrated and their data encrypted and possibly stolen as well [13]. Here the tool builders provide almost everything needed to carry out the crime: the malware, dashboards to monitor progress in compromising systems, and storage systems to hold exfiltrated data (and publish it should negotiations fail). All that remains for the buyer of this system to do is to deliver the malware, set the price the victim must pay and collect (and subsequently launder) the money.

Overall, when one looks at tools and infrastructure there is limited sophistication to be seen here. There is of course the ability to hide away the details of making an exploit actually work so that poorly skilled tool purchasers can be successful, or making it easy for a kit user to understand how to deploy their new acquisition. There is some innovation, mainly in combining known techniques together, but 'phishing kits', for example, have barely changed in 15 years [96]. Code is rarely written from scratch, with exploits mostly stolen from public dumps on websites like pastebin, shared on websites (as Mirai was [7]) or bought directly. Equally, as far as can be determined, infrastructure providers rarely appear to develop any sophisticated understanding of the systems they run or the environment in which they operate, working instead on trial-and-error basis.

However, while there is little technical sophistication in this part of the cybercrime ecosystem, that is not to say that there is no ingenuity or innovation. Where tool providers are concerned, ingenuity tends to be focused on smooth and flashy UX and marketing – dressing up trivial changes to payloads as a new attack method, organising flash sales and improving sales patter on forums, or adding skull symbols and other attractive graphics to a command line tool [41].

For infrastructure providers, the goal is to keep systems stable and usable, so there is some innovation and experimentation in the supportive techniques of system administration, moderation, customer service, and avoiding detection by law enforcement, payment services, and ISPs. These innovations in administrative work can be genuinely creative – setting up automated transfers between a network of PayPal accounts,

experimenting with hosting bots in different countries, etc. – however they do not rely on deep technical knowledge. Often the deviant admin will experiment with a range of different approaches at random until they find one that works, without ever knowing why. The highest level of sophistication generally observed in this kind of work is in fact in *automation* – few of the infrastructure providers in the volume cybercrime-as-a-service ecosystem have even the limited capabilities necessary to 'automate the boring stuff', so that those who can automate will gain a huge advantage, often leasing out infrastructure to other groups and capturing their less-skilled rivals (who need to put in enormous manual effort to keep their systems running at all) into affiliate schemes [23].

4.4 Entrepreneurs

The users of the tools and infrastructure we have just described form a separate group. They may understand how the tools that they use actually work, but there is no need for them to do so. They are, as Anderson et al. [6] set out, best viewed as entrepreneurs, and their creativity is often social – they rarely have any real knowledge of technology, and their success comes more from successfully delivering the exploit or incorporating it into a business, scam, or crime script. Interestingly, most of the discussions which are self-described as 'advanced hacking' on cybercrime forums, such as those collected in the CrimeBB dataset [72] fall into this category, despite the lack of any real technical knowledge, let alone sophistication.

What is seen is the sharing of criminal scheme-patterns or automation scripts, or discussions and training materials in the basic mechanics of tool use. Entrepreneurs cover the vast majority of the actors which cause direct, serious harm in the cybercrime ecosystem. They take the services, skills, and tools provided by other actors and groups and put them into practice, generating profits and causing harm to victims. The groups they form to operate their 'business' may well incorporate one or two more technical members, however technical ingenuity and skill are rarely the determinants of success. This formation can be seen across a wide range of cybercrimes – denial of service, ransomware gangs, romance scammers, and fraudsters. Where a group becomes particularly successful, there is occasionally movement into the provision of illicit infrastructure, as the crime script is semi-automated, turned into a service, and sold to others through affiliation schemes [73].

Operational security practices (OPSEC) are particularly important to this group, as they are the most likely to engage in direct engagement with victims at scale, and hence the most likely to come to the reactive attentions of law enforcement. While technical sophistication is indeed possible in the OPSEC domain, these more sophisticated practices are generally confined to the skilled actors we described in Section 4.3 above. However, criminal entrepreneurs do use basic security practices such as accessing systems over VPNs or using encrypted messaging services, as these require very little knowledge to employ. Using these effectively to evade law enforcement is another matter (and these actors generally overestimate their own OPSEC abilities [53]), however security economics operates in these actors' favour, and if they can increase the money and effort which needs to be expended in order to identify them, fairly simple measures will often be sufficiently protective – unless law enforcement decides to specifically concentrate on identifying them.

There is a great deal of creativity shown by this group – but the skills shown are entrepreneurial rather than technical. Generally, we see business model innovation, creativity in creating crime scripts, advertising and marketing gimmicks, and hands-on experimentation with getting systems to work in practice. Deep knowledge of how things work isn't as important as persistence and messing around – in many ways this is 'folk magic' rather than true wizardry [23].

The economics of criminal entrepreneurship are fairly well-understood [50]. These generally involve selection of suitable targets and the development of lucrative business models. Skill barriers are important in directing activity, with entrepreneurial efforts generally moving to forms of crime where the skill level required to participate has been gradually worn down over time; through improvements in the usability of tools, supportive services, and the provision of infrastructure. This allows more developed business models and crime scripts to emerge and be shared socially within cybercrime communities.

Sophistication and usability also drive practices on the 'victim' side, as many frauds involve at least some action on the part of the target, who cannot be assumed to have any technical skills whatsoever. Despite the apparent relative ease of e.g. installing Tor and accessing an Onion Service, the targets of deviant entrepreneurs have historically struggled with even the most basic use of, for example, Bitcoin to transfer funds to pay a ransom, and so usability and simplicity drives this part of the ecosystem as well. This can be seen both with the proliferation of supportive technical services (such as the brief use of Tor2Web to allow ransomware victims to access an onion service without downloading Tor) and even full customer service operations to facilitate the transfer of cryptocurrency. Thus, even very minor usability or convenience barriers can pose serious challenges to entrepreneurial efforts – evidenced by the preponderance of ample cybercrime and drug buying services on the 'clearnet', which far outnumber their 'darknet' equivalents.

4.5 'Script kiddies' and petty fraudsters

The final truly criminal grouping are disparaged as 'script kiddies' (or 'skids'). These are generally at the beginning of their cybercriminal careers and are just about capable of using a cybercrime service or following a 'how to' guide. There is occasionally some innovation here, but of a non-technical form as this group incorporates the tools that they come across into 'crime scripts'. This is often the preserve of 'beginner hacking' sections on forums – not hacking at all, but rather the operation of romance scams, other petty frauds, and 'get rich quick' schemes like drop-shipping.

When members of this group do come to the attention of Law Enforcement, perhaps because they are traced through their purchases of booter services or hacking tools, they may not end up in court – instead they will often get a formal warning or be sent on a diversionary programme that will attempt to interest them in a legitimate career in computer security. Alternatively, their purchase of a surveillance tool may be an incidental part of an otherwise entirely non-technical crime such as domestic abuse, harassment, voyeurism or fraud [89, 43].

4.6 Hacktivists and other non-monetary motives

Some of the people who attack computer systems are not doing it to make money for themselves but for 'higher motives'. They may consider themselves to be part of the security industry and are providing demonstrations of insecurities so that fixing them becomes a higher priority. They may have political motives and either wish to knock systems offline, or steal material, in both cases because they believe the publicity will advance their cause [80, 65]. This group of people cause some anguish to criminologists because they are present in the ecosystem but standard theories about criminal careers clearly won't apply to them. As a result they probably get more attention in the literature than they deserve – they are rarely seen, and most instances of hacktivism directly relate to real-world conflicts ('patriotic hackers' attacking India or Pakistan, Israel or Palestine, Russia or Ukraine...).

Where there are more personal motives for law-breaking, for example, Gary McKinnon sought evidence of coverups of UFO activity in US military and NASA computers, the criminal justice has some difficulty dealing with the cases. Quite clearly McKinnon broke US (and UK) computer crime laws, but the extradition process engaged in by the US, and the maximum penalties they threatened, did not give the impression of a proportionate response. Eventually the UK Home Secretary took the decision to override the courts and refuse the US request [8, 90]. The absence of a financial motive makes existing diversionary programmes almost irrelevant and the rarity of such cases means that new approaches are unlikely to be seen any time soon.

Hacktivists are few in number, may act totally alone or in small, medium or large groups, and their technical approaches vary enormously. No general assessment of their sophistication is possible.

5 Criminology theory and cybercrime sophistication

There are broader consequences of misleading depictions of cybercrime as invariably being sophisticated. In particular, we argue that this is leading to the mis-specification of theoretical models of cybercrime, which can be influential in directing police activities and priorities. Someone who makes tools for others develops different skills from someone who uses these tools in a complex setting. Someone who pays this person to use these tools for them may also display a range of skilled and complex behaviours – but in the domain of social organisation, incorporating these technologies into complex scams. Each of these different forms of activity presents different risk, motivations, and pathways, and needs to be dealt with differently.

The mis-specified model generally focuses around a hypothecated increasing sophistication over time for those involved in cybercrime, borrowing some elements of a social learning model of crime [2], and resonates with the 'funnel' or 'pathway' model seen in some radicalisation scholarship, in which involvement and linked factors progressively deepen over time. Although this model underpins substantial law enforcement activity, such as the UK's PREVENT programme, it is increasingly criticised by academic researchers [52]. In a similar model, involvement in cybercrime is seen in classic criminological terms as supported by the slow internalisation of community norms, practices, and skills, where the individual becomes less wary of committing deviant acts, progresses to more serious offences, and develops proficiency in the argot and practices of the subculture or community [40].

While this model has much to recommend it as a theory of deviant social organisation, in the case of cybercrime it has generally been bound up with the idea that a core part of this pathway is the development of substantial technical skill, sophistication, and understanding of computers and the Internet, which allows individuals to progress up the 'ladder' towards elite status. However, more recent research suggests strongly that an ecosystem model, defined by specialised niches and forms of work that interrelate, rather than a linear progression model of low to high sophistication, is more appropriate. This is backed up by a growing body of evidence for the industrialisation of the cybercrime ecosystem – rather than communities with higher- and lower-skilled members on a continuum, cybercrime is defined by a scrappy association of scams, affiliate schemes, tools, and services, in which success may be determined as much by good customer service, luck in evading law enforcement, or flashy graphics and marketing as it is by technical sophistication [41, 67].

Thus, those involved in cybercrime generally appear to learn within niches rather than progressing to 'elite' status, often flitting between a range of different crime types at similar levels of sophistication [56]. While the progress of individuals may well be usefully described as a pathway, it is crucial to recognise that very few of the 'pathways' of cybercrime lead to either the sophisticated technical skills needed for exploit development, or the project management, attack frameworks, coding, and other skills needed for most jobs in the cybersecurity industry [23].

Technical sophistication (or lack thereof) is also crucial in shaping the *economic* factors at play in cybercrime. Within illicit economies, engagement in various activities is determined not only by rational risk-reward calculations, but by a wide range of other factors, including the availability of the technical skills and practices needed to do so. In cybercrime, lack of technical skill does not pose a direct *barrier* to engaging in particular forms of crime – although the basic abilities needed to operate tools and pull off scams successfully do need to be learned over time, they are not necessarily technically sophisticated. This gives rise to something of a skill economy, in which a small number of more skilled practitioners are able to support wider illicit activity.

The incorporation of technologies and infrastructures into this skill economy is itself one of the central reasons that cybercrime need not be a particularly skilled pursuit. One of the transformative capacities of digital technologies lies in their ability to be invoked or 'conjured' by naïve users – such that a crime involving a number of apparently-technical capabilities may in fact simply involve following a YouTube tutorial to copy and past code into a webpage, with little understanding of how this will take advantage of a well-established facile vulnerability [38].

This putative 'criminal' may show very little commitment to deviance, have minimal appreciation or intention of harm, demonstrate trivial technical abilities or understanding, and yet have unwittingly completed the very same steps that a GCHQ-trained professional penetration tester might have taken to compromise a poorly-defended web service. Equally, it only requires a single skilled actor to create a useful tool or release the code for an exploit on a forum to reduce dramatically the skill barrier for a lucrative form of crime – as we have seen on numerous occasions, such as the release of the Mirai source code on HackForums [7].

Where business models innovate to operate these tools as services or embed them in infrastructures, the skill barrier reduces yet further – as seen with the Zeus-as-a-service industry [42] or the spread of DDoS services [44]. This means that the skill barrier to cybercrime is nonlinear – if one person can overcome it, they can generally reduce it for others. This means that cybercrime economies, as we explored above, show relatively few links between the different roles and niches within the wider ecosystem. There is little relationship between the skills and pathways of the petty fraudsters who make up much cybercrime and the more complex work of vulnerability research or enterprise system compromise. In other words, despite some exceptions, these are largely interrelated but separate pathways and domains of activity, rather than a continuum.

For judges, activists, police, defenders, and policymakers (and the individuals in these cybercrime communities themselves) this poses serious problems for the administration of justice – as directly demonstrated in the wildly varying punishments meted out for cybercrime, which range from draconian Internet banning orders, threats of extradition and the prospect of decades in jail for trivial experimentation and curiosity; to diversion-based disposals for involvement in complex, high-harm frauds and illicit enterprise.

Further, ideas of technical sophistication themselves play an important cultural role in cybercrime communities and in the individual journeys taken by those who get involved in them [37, 38]. Advanced technical mastery has historically been a core value of the underground 'hacker culture' and so plays an important role in organising activities of the subculture. Increasingly, as cybercrime has become industrialised, this has given way to a more entrepreneur-focused set of cultural values, with the 'hacker ethic' decreasingly important in many of the communities which make up the cybercrime underground. Despite this, we argue that media and police narratives play an important part in *legitimising* the core values of the illicit hacker subculture – this persistent image of cool, deviant geniuses remains an attractive one. Thus, depicting cybercrime as high damage, technically sophisticated, and lucrative may well unintentionally be doing the marketing of these groups for them, and glamorising what is in fact a rather boring and low-skilled business.

6 Discussion

We have discussed in this paper the major different components of the cybercrime ecosystem and the distinct practices, skills, and levels of technical sophistication in each. Importantly, we have focused this around *criminalized* activities rather than 'hacking' per se, on which much of the literature dwells, and which, we argue, is a decreasingly useful concept for explaining or understanding most of the work and people involved in cybercrime. While cultures of technological experimentation and innovation were, in the 80's, 90's, and 2000's, likely more closely linked to the cybercrime underground, and in turn, this underground had more viable paths to the legitimate security industry [32], we argue that in an era of industrialised cybercrime this is now the case only for a small minority of individuals. The main conclusion of this work will be of little surprise to most in industry – the vast majority of people involved in cybercrime never develop any significant technical sophistication, nor is there any particular need for them to do so in order to continue to make money, cause harm, and compromise systems.

This has a direct relevance for criminal justice policy. If many of the people identified or caught by law enforcement are in fact purchasing services or simply deviant entrepreneurs then few of them are likely to have the professional skills or technical interests which might suit a career in cybersecurity. The efforts

to redirect them into other careers might best be focused on the far smaller number of tool-builders and infrastructure providers. There is already some concern that PREVENT-style programmes which assume a linear pathway and engage with a wide 'pre-criminal' population in fact simply engage in net widening rather than actually diverting people away from the criminal justice system. The relatively low harm engaged in by those contacted by PREVENT-style programmes generally serves to reinforce the perceived severity and intractability of those who are found to have committed criminal offences, thus, the stigmatising reach of criminal justice is merely extended to a wider population, and so there is little overall reduction in jail sentences given. There needs to be further consideration given to the possibilities of negative unintended consequences, for example, of stigmatising those with ASD, contributing to self-fulfilling 'labelling' of young people as deviant, and more broadly to the problem of separating simple curiosity from crime.

A similar fallacy has been influential in counter-radicalisation policy – i.e. that individuals follow an escalator-style pathway of increasingly extreme beliefs, which are perceived as the main drivers of eventual violent action. This has been widely critiqued in recent years [52]. For cybercrime, an analogous narrative has taken hold: of individuals following an escalator of increasing *technical sophistication and skill*, which is hypothecated as the main driver of eventual serious crime. In reality, many of the most serious and harmful cybercrimes involve the lowest levels of technical skill – especially where they are combined with domestic violence, harassment, stalking, and other gendered forms of harm.

Understanding sophistication in the cybercrime ecosystem is also important for what works in terms of interventions. The UK's National Crime Agency have had demonstrable successes with strategic communications adverts which assume very low skill and knowledge, targeting the least skilled people in the ecosystem [24]. Arresting individual key players generally appears less successful, as the heads of these crime businesses are entrepreneurs, and the skilled work is being done elsewhere. Conversely, where a cybercrime economy loses the few people able to maintain the infrastructure, either through police action or ageing out, the limited number of empirical studies performed thus far show a substantial disruption [25].

We further draw on this systematic analysis of cybercrime ecosystems to argue that 'cyber enabled' and 'cyber dependent' are increasingly insufficient terms with which to describe the reality of contemporary cybercrime [95]. These terms, distinguishing traditional crimes with an online or technological component (such as fraud) from those which could not exist without networked computer systems (such as intrusion-based hacking or denial of service) have been extremely useful as a way for the police to categorise work so that the right people work on the right things. However, for the purposes of understanding cybercrime, these now serve neither to distinguish the levels of technical sophistication involved in a crime (with some romance frauds utilising deepfake machine learning tools and some denial of service attacks operable with no technical experience whatsoever), nor to meaningfully distinguish separate fields of illegal activity. We observe similar (and sometimes the same) people, infrastructures, kinds of work, and even levels of technical sophistication involved in both cyber-'enabled' and -'dependent' crimes.

7 Related work

We are far from alone in calling out the almost ubiquitous use of 'sophisticated' to describe successful cyberattacks. In 2020 Kleinman observed that the companies announcing data breaches inevitably made the three points "We were the victim of a sophisticated attack; We take our customers' privacy seriously; to date, we have no evidence that passwords were accessed or data has been used". He goes on to observe "In almost every supposed 'sophisticated' attack, well-known and previously identified methods and vulnerabilities are the sources of exploitation" and he questions whether the companies had actually been serious about privacy [48].

The rise of cybercrime-as-a-service, and the broader industrialisation, specialisation, and professionalisation of the cybercrime ecosystem is now a well-established phenomenon, and has been noted and analysed by practitioners in the cybersecurity industry for some time. The economics of the vulnerability markets

associated with these tools and attacks have also been the subject of substantial research [4]. Interconnection and chain effects within the different parts of cybercrime ecosystems have also been the subject of previous scholarship [74]. Possibly the earliest example of the move from selling tools to running full-scale professionalised services occurred with booter services, as people looked for ways to monetise their botnets and exploitation of amplification/reflection vulnerabilities became more common [80]. Subsequently, ransomware has similarly evolved to offer a service model [13], and these entrepreneurial business models have also spread into other forms of online crime – with digital service infrastructures emerging for a range of purposes, including social media bots for online harassment and misinformation, spyware-as-a-service, and similar. Industry perspectives include in-depth investigations of the dynamics within these gangs, e.g. those carried out by analysts such as the Grugq (and others), who has documented the dynamics of professionalised 'CrimeOps' and the relative unimportance of technical sophistication for a successful cybercrime group [34].

Our view that cybercrime is seldom much to do with sophisticated paradigm changing attacks constructed by hacking deities finds an echo in "The Stamos Hierarchy of the Actual Bad Stuff that Happens Online to Real People" which Alex Stamos (ex CSO of Yahoo and Facebook) introduced in a 2019 USENIX Security keynote. He explained that most harm on the Internet was "abuse", tools being used as they were meant to be used, but with ill-intent. The, relatively tiny, "InfoSec" top-most portion of his illustrative triangle covered all password, patching and mis-configuration issues that have been known for years and still caused significant security issues. Only the single "pixel" at the top of his triangle related to 0-days and, only some small sub-unit of that pixel, he argued, related to the topics, such as side-channels, addressed in the USENIX programme [84].

Our observations on 'APT' chime with those of NCSC technical director Ian Levy who has described them as 'Adequate Pernicious Toe-rags'. He argued that the sophistication of these groups had been overstated by threat intelligence companies with a vested interest in conjuring up sophisticated attackers to make their own services, which Levy describes as effectively 'magic amulets', seem more advanced [57].

Ohm argued, 15 years ago, that policymakers passed overbroad, ambiguous cybercrime laws intended to ensnare the 'Superuser' – "a mythic figure: difficult to find, immune to technological constraints, and aware of legal loopholes". He argued that "fear of the Internet" was at the root of these policy failures and that "computer experts are as susceptible as laypeople to exaggerate the power of the Superuser" [69].

The classic cultural image of the *hacker* still has a strong hold on policy, practice, and research on cybercrime [86]. This puts an emphasis on a very specific form of advanced technical sophistication as being at the heart of 'hacking', including experimentation, excitement, creativity, and hard technical mastery [33]. Criminologists have long realised that there is considerable variation between the abilities of different cybercriminals. Much of the early research literature on cybercrime explicitly addressed sophistication as a central taxonomic or explanatory factor [20, 59, 94]. The consensus view, a 'trickle down' theory, has become one of a relatively small group of highly technically proficient criminals who build tools for other less proficient criminals to use. Hutchings provides a good overview of various versions of this model [40], which is clearly related to the one we have presented. However, we are more sceptical about the extent to which the most cutting edge techniques become widely adopted, we feel that it is useful to distinguish the entrepreneurs and infrastructure providers from the script kiddies, we argue that even most tool builders rarely display genuine sophistication, and we put hacktivists, of whatever skill, into a category of their own. While, as documented in an extremely valuable recent report by Goerzen and Coleman, in prior years it was more common for those in the hacker underground to move into professionalised security careers [32], we argue that the contemporary ecosystem, in which volume frauds and illicit services predominate, decreasingly facilitates these kinds of transitions.

As the cybercrime economy has begun to industrialise and achieve genuine scale, criminologists have generally moved away from purely subcultural framings, or the 'serious organised crime' model, instead understanding the ecosystem as defined by mass waves of entrepreneurial experimentation with business models, crime scripts, and marketing approaches, with advances in the underlying technologies and vulnerabilities occurring only rarely [30, 18, 73, 56].

8 Conclusions

We started this paper by giving several examples of some fairly humdrum cybercrimes being described as sophisticated. Although some people will be unaware of how straightforward most cybercrime is ("any sufficiently advanced technology is indistinguishable from magic") we believe the main reason for inflating the prowess of the cybercriminals is that there is an alignment of incentives, with everyone involved having clear incentives to talk up the sophistication of what has occurred.

We then considered how correctly identifying sophistication matters. Within the criminal justice system sophisticated crimes are treated more seriously, perhaps thereby providing a disincentive for the criminals themselves to claim that their actions were sophisticated. More broadly we explained how the perception of ubiquitous sophistication adversely affected policy initiatives.

We then described six groups of cybercrime actors and set out how sophisticated we believe their activities actually to be. We indicated that these groups are linked together rather less than mainstream criminology theory currently suggests. In particular we do not believe that the cutting edge discoveries of novel techniques have much immediate influence on cybercrime and we treat hacktivists as being rather separate – partly because their skillset varies very widely indeed and partly because we believe they are relatively few in number and, at present, they have little real impact beyond distracting attention from genuinely harmful and criminal activities.

As we have demonstrated, these distinct components of the cybercrime ecosystem not only involve very different forms of sophistication and sociotechnical practice, but operate along diverging economic and social lines, driven by their own motivations and incentives. It is important to define what kinds of creativity, sophistication, and technical ability are actually present at each of these levels. Rather different forms of sophistication are present at each of the levels we identify within the cybercrime ecosystem. Some aspects involve genuine technical knowledge, others professional skills and practices, still others system administration and customer service skills, however these account for a relatively small group of actors. Although a handful of genuinely skilled actors have an outsize influence on the rest of the cybercrime ecosystem, most of the skills involved are focused on entrepreneurship rather than IT literacy or engineering creativity.

We argue that this mis-attribution of sophistication is not a mere irritation or fringe issue, i.e. something which security professionals and academics might decry on social media but which bears no real impact; it is in fact contributing to a far wider mood music within law enforcement and sentencing practice, and academia that has had serious negative effects. As long as this narrative of sophistication persists, it creates perverse incentives to widen the net of law enforcement, to design interventions along misleading lines, and to draw attention away from the real drivers of harm.

The fear of cybercrime as a problem of advanced technology and sophisticated, skilled actors, rather than of creative and deviant entrepreneurship in contemporary digital societies, contributes to policy that sees it a technical problem with technical solutions, and focuses the wrong kinds of interventions at the wrong kinds of actor. In our view, what is desperately needed is a more sophisticated approach.

References

[1] Sadia Afroz, Vaibhav Garg, Damon McCoy, and Rachel Greenstadt. Honor among thieves: A common's analysis of cybercrime economies. In *APWG eCrime Researchers Summit*, pages 1–11. IEEE, 2013.

[2] Ronald L Akers. *Social learning and social structure: A general theory of crime and deviance*. Routledge, 2017.

[3] Luca Allodi. The heavy tails of vulnerability exploitation. In *International Symposium on Engineering Secure Software and Systems*, pages 133–148. Springer, 2015.

[4] Luca Allodi. Economic factors of vulnerability trade and exploitation. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*, pages 1483–1499, 2017.

[5] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Carlos Gañán, Tom Grasso, Michael Levi, Tyler Moore, and Marie Vasek. Measuring the changing cost of cybercrime. In *Workshop on the Economics of Information Security*, 2019.

[6] Ross Anderson, Richard Clayton, Rainer Böhme, and Ben Collier. Silicon den: Cybercrime is entrepreneurship. In *Workshop on the Economics of Information Security*, 2021.

[7] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, Alex J Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. Understanding the Mirai botnet. *USENIX Security Symposium*, pages 1093–1110, 2017.

[8] Paul Arnell and Alan Reid. Hackers beware: the cautionary story of Gary McKinnon. *Information & Communications Technology Law*, 18(1):1–12, 2009.

[9] Ashish Arora, Ramayya Krishnan, Rahul Telang, and Yubao Yang. Impact of vulnerability disclosure and patch availability – an empirical analysis. In *Workshop on the Economics of Information Security*, 2004.

[10] Ashish Arora, Rahul Telang, and Hao Xu. Optimal policy for software vulnerability disclosure. *Management Science*, 54:642–656, 04 2008.

[11] W Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394):116–131, 1989.

[12] Miriam H Baer. Unsophisticated sentencing. *Wayne L. Rev.*, 61:61, 2015.

[13] Kurt Baker. Ransomware as a service (RaaS) explained, 2022. URL: https://www.crowdstrike.com/cybersecurity-101/ransomware/ransomware-as-a-service-raas/.

[14] Gabriel Bassett, C. David Hylender, Philippe Langlois, Alexandre Pinto, and Suzanne Widup. Verizon 2021 Data Breach Investigations Report, 2022. URL: https://verizon.com/dbir/.

[15] BBC. General election 2019: Labour Party hit by second cyber-attack, 2019. URL: https://www.bbc.co.uk/news/election-2019-50388879.

[16] Ian Beer and Samuel Groß. A deep dive into an NSO zero-click iMessage exploit: Remote code execution. https://googleprojectzero.blogspot.com/2021/12/a-deep-dive-into-nso-zero-click.html, 2021.

[17] Frank O Bowman III. Damp squib: The disappointing denouement of the sentencing commission's economic crime project (and what they should do now). *Federal Sentencing Reporter*, 27(5):270–283, 2015.

[18] Roderic Broadhurst, Peter Grabosky, Mamoun Alazab, Brigitte Bouhours, and Steve Chon. An analysis of the nature of groups engaged in cyber crime. *International Journal of Cyber Criminology*, 8(1):1–20, 2014.

[19] John Chambers. What does the Internet of Everything mean for security?, 2015. URL: https://www.weforum.org/agenda/2015/01/companies-fighting-cyber-crime/.

[20] Kim-Kwang Raymond Choo. Organised crime groups in cyberspace: a typology. *Trends in organized crime*, 11(3):270–295, 2008.

[21] Arthur C. Clarke. Clarke's Third Law on UFO's. *Science*, 159(3812), 1968.

[22] Stanley Cohen. *Visions of social control*. Polity Press Cambridge, 1985.

[23] Ben Collier, Richard Clayton, Alice Hutchings, and Daniel Thomas. Cybercrime is (often) boring: maintaining the infrastructure of cybercrime economies. In *Workshop on the Economics of Information Security*, 2020.

[24] Ben Collier, Daniel R Thomas, Richard Clayton, and Alice Hutchings. Booting the booters: Evaluating the effects of police interventions in the market for denial-of-service attacks. In *Proceedings of the Internet Measurement Conference (IMC)*, pages 50–64. ACM, 2019.

[25] Ben Collier, Daniel R Thomas, Richard Clayton, Alice Hutchings, and Yi Ting Chua. Influence, infrastructure, and recentering cybercrime policing: evaluating emerging approaches to online law enforcement through a market for cybercrime services. *Policing and Society*, pages 1–22, 2021.

[26] Marco Cova, Christopher Kruegel, and Giovanni Vigna. There is no free phish: An analysis of "free" and live phishing kits. In *2nd USENIX Workshop on Offensive Technologies (WOOT 08)*, San Jose, CA, 2008. USENIX Association.

[27] Crown Prosecution Service. Legal Guidance: Computer Misuse Act, 2020. URL: `https://www.cps.gov.uk/legal-guidance/computer-misuse-act`.

[28] Giovanni Dosi. Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change. *Research policy*, 11(3):147–162, 1982.

[29] Joseph Farrell and Garth Saloner. Standardization, compatibility, and innovation. *the RAND Journal of Economics*, pages 70–83, 1985.

[30] Vaibhav Garg and L Jean Camp. Why cybercrime? *Acm Sigcas Computers and Society*, 45(2):20–28, 2015.

[31] Frank W Geels. Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. *Research Policy*, 31(8-9):1257–1274, 2002.

[32] Matt Goerzen and Gabriella Coleman. Wearing many hats, the rise of the professional security hacker. *Data and Society*, 2022.

[33] Andrew Goldsmith and David S Wall. The seductions of cybercrime: Adolescence and the thrills of digital transgression. *European Journal of Criminology*, 2019. `doi:10.1177/1477370819887305`.

[34] the Grugq. Crimeops: The operational art of cyber crime, 2020. URL: `https://sec.okta.com/articles/2020/08/crimeops-operational-art-cyber-crime`.

[35] Hermann Härtig, Claude-Joachim Hamann, and Michael Roitzsch. The mathematics of obscurity: On the trustworthiness of open source. In *Workshop on the Economics of Information Security*, 2010.

[36] Alex Hern. TalkTalk hit with record £400k fine over cyber-attack. *The Guardian*, 2016. URL: `https://www.theguardian.com/business/2016/oct/05/talktalk-hit-with-record-400k-fine-over-cyber-attack`.

[37] Thomas J Holt. Subcultural evolution? examining the influence of on-and off-line experiences on deviant subcultures. *Deviant Behavior*, 28(2):171–198, 2007.

[38] Thomas J Holt. Examining the role of technology in the formation of deviant subcultures. *Social Science Computer Review*, 28(4):466–481, 2010.

[39] John Howells. The response of old technology incumbents to technological competition – does the sailing ship effect exist? *Journal of Management Studies*, 39(7):887–906, 2002.

[40] Alice Hutchings. Crime from the keyboard: organised cybercrime, co-offending, initiation and knowledge transmission. *Crime, Law and Social Change*, 62(1):1–20, 2014.

[41] Alice Hutchings and Richard Clayton. Exploring the provision of online booter services. *Deviant Behavior*, 37(10):1163–1178, 2016.

[42] Alice Hutchings and Richard Clayton. Configuring Zeus: A case study of online crime target selection and knowledge transmission. In *Proceedings of the 2017 APWG Symposium on Electronic Crime Research (eCrime)*, pages 33–40. IEEE, 2017.

[43] Alice Hutchings and Sergio Pastrana. Understanding ewhoring. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 201–214. IEEE, 2019.

[44] Mohammad Karami and Damon McCoy. Rent to pwn: Analyzing commodity booter DDoS services. *USENIX ;login*, 38(6):20–23, 2013.

[45] Neal Kumar Katyal. Criminal law in cyberspace. *University of Pennsylvania Law Review*, 149(4):1003–1114, 2001.

[46] David K Kessler. Cybercrime sentencing. *Dep't of Just. J. Fed. L. & Prac.*, 69:107, 2021.

[47] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors. In *Proceedings – International Symposium on Computer Architecture*, 06 2014.

[48] Leonard Kleinman. Cyberattacks: Just how sophisticated have they become? *Forbes Technology Council, Council Post*, 2020. URL: https://www.forbes.com/sites/forbestechcouncil/2020/11/03/cyberattacks-just-how-sophisticated-have-they-become/.

[49] Vadim Kotov and Fabio Massacci. Anatomy of exploit kits. In Jan Jürjens, Benjamin Livshits, and Riccardo Scandariato, editors, *Engineering Secure Software and Systems*, pages 181–196, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

[50] Nir Kshetri. The simple economics of cybercrimes. *IEEE Security & Privacy*, 4(1):33–39, 2006.

[51] Jeanne Kuang. Missouri prosecutor declines to charge St. Louis Post-Dispatch reporter Parson targeted, 2022. URL: https://www.kansascity.com/news/politics-government/article258315738.html.

[52] Arun Kundnani. Radicalisation: the journey of a concept. *Race & class*, 54(2):3–25, 2012.

[53] Isak Ladegaard. We know where you are, what you are doing and we will catch you: Testing deterrence theory in digital drug markets. *The British Journal of Criminology*, 58(2):414–433, 2018.

[54] Ralph Langner. Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security & Privacy*, 9(3):49–51, 2011.

[55] Rebecca Ledingham and Richard Mills. A preliminary study of autism and cybercrime in the context of international law enforcement. *Advances in Autism*, 2015.

[56] Eric Rutger Leukfeldt and Thomas J Holt. Cybercrime on the menu? examining cafeteria-style offending among financially motivated cybercriminals. *Computers in Human Behavior*, 126:106979, 2022.

[57] Ian Levy. National scale cyber security. In *USENIX Enigma*, 2017. URL: https://www.usenix.org/conference/enigma2017/conference-program/presentation/levy.

[58] Alliyza Lim, Neil Brewer, and Robyn L Young. Revisiting the relationship between cybercrime, autistic traits, and autism. *Journal of Autism and Developmental Disorders*, pages 1–12, 2021.

[59] Howard F Lipson. Tracking and tracing cyber-attacks: Technical challenges and global policy issues. Technical report, Carnegie-Mellon Univ Pittsburgh Pa Software Engineering Inst, 2002.

[60] Log4j team. Fixed in Log4j 2.16.0 (Java 8) and Log4j 2.12.2 (Java 7), 2020. URL: https://logging.apache.org/log4j/2.x/security.html.

[61] Dhia Mahjoub. Behaviors and patterns of bulletproof and anonymous hosting providers. In *USENIX Enigma*, 2017. URL: https://www.usenix.org/conference/enigma2017/conference-program/presentation/mahjoub.

[62] Derek Manky. Cybercrime as a service: a very modern business. *Computer Fraud & Security*, 2013(6):9–13, 2013.

[63] Lesley McAra and Susan McVie. Youth crime and justice: Key messages from the Edinburgh study of youth transitions and crime. *Criminology & Criminal Justice*, 10(2):179–209, 2010.

[64] Merriam-Webster.com. Dictionary, 2022. URL: https://www.merriam-webster.com/dictionary/sophisticated.

[65] Stefania Milan. Hacktivism as a radical media practice, 2015. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2901929.

[66] Tyler Moore and Richard Clayton. How hard can it be to measure phishing? In *Mapping and Measuring Cybercrime, Oxford, UK*, 2010.

[67] Roberto Musotto and David S Wall. More Amazon than Mafia: analysing a DDoS stresser service as organised cybercrime. *Trends in Organized Crime*, pages 1–19, 2020.

[68] Jens Neisius and Richard Clayton. Orchestrated crime: The high yield investment fraud ecosystem. In *2014 APWG Symposium on Electronic Crime Research*, pages 48–58. IEEE, 2014.

[69] Paul Ohm. The myth of the superuser: Fear, risk, and harm online. *UC Davis Law Review*, 41:1327–1402, 2008.

[70] OpenSSL. TLS heartbeat read overrun (CVE-2014-0160). URL: https://www.openssl.org/news/secadv/20140407.txt.

[71] Rebekah Overdorf, Marc Juarez, Gunes Acar, Rachel Greenstadt, and Claudia Diaz. How unique is your .onion? an analysis of the fingerprintability of Tor onion services. In *CCS '17*, pages 2021–2036, New York, NY, USA, 2017. Association for Computing Machinery.

[72] Sergio Pastrana, Daniel R Thomas, Alice Hutchings, and Richard Clayton. CrimeBB: Enabling cybercrime research on underground forums at scale. In *Proceedings of the 2018 World Wide Web Conference*, pages 1845–1854, 2018.

[73] Maria Grazia Porcedda and David S Wall. Cascade and chain effects in big data cybercrime: Lessons from the TalkTalk hack. In *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 443–452. IEEE, 2019.

[74] Maria Grazia Porcedda and David S Wall. Modelling the cybercrime cascade effect in data crime. In *2021 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 161–177. IEEE, 2021.

[75] Sam Ransbotham. An empirical analysis of exploitation attempts based on vulnerabilities in open source software. In *Workshop on the Economics of Information Security*, 2010.

[76] Sam Ransbotham and Sabyasachi Mitra. The impact of immediate disclosure on attack diffusion and volume. In Bruce Schneier, editor, *Economics of Information Security and Privacy III*, pages 1–12, New York, NY, 2013. Springer New York.

[77] Shaked Reiner. Golden SAML revisited: The Solorigate connection, 2020. URL: https://www.cyberark.com/resources/threat-research-blog/golden-saml-revisited-the-solorigate-connection.

[78] Eric Rescorla. Is finding security holes a good idea? *IEEE Security & Privacy*, 3(1):14–19, 2005.

[79] Jamie Saunders. Tackling cybercrime – the UK response. *Journal of cyber policy*, 2(1):4–15, 2017.

[80] Molly Sauter. 'LOIC will tear us apart': The impact of tool design and media portrayals in the success of activist DDoS attacks. *American Behavioral Scientist*, 57(7):983–1007, 2013.

[81] John Simpson and Edmund Weiner (eds). *Oxford English Dictionary, Second Edition*. Oxford University Press, 1989.

[82] Spalding Voice. Local schoolboy (17) ordered to pay back over £2.1m after internet fraud. https://www.spaldingvoice.co.uk/local-schoolboy-17-ordered-to-pay-back-over-2-1m-after-internet-fraud/, 2021.

[83] Matt Spencer. Creative malfunction: Finding fault with rowhammer. *Computational Culture*, 8, 2021.

[84] Alex Stamos. Keynote: Tackling the trust and safety crisis. In *USENIX Security '19*, 2019.

[85] Kevin F Steinmetz. Executing effective social engineering penetration tests: A qualitative analysis. *Journal of Applied Security Research*, pages 1–21, 2021.

[86] Kevin F Steinmetz, Thomas J Holt, and Karen M Holt. Decoding the binary: Reconsidering the hacker subculture through a gendered lens. *Deviant Behavior*, 41(8):936–948, 2020.

[87] Jack Stubbs. Exclusive: UK's Labour sticks to 'basic' $20 cyber defense after attacks, emails show, 2019. URL: https://www.reuters.com/article/us-britain-election-cyber-exclusive-idCAKBN1XT27R.

[88] TalkTalk. Website attack affecting our customers. https://web.archive.org/web/20151022211310/http://help2.talktalk.co.uk/oct22incident, 2015.

[89] Leonie Tanczer, Isabel Lopez Neira, Simon Parkin, Trupti Patel, and George Danezis. Gender and IoT research report. *University College London, white paper*, 2018.

[90] Alan Travis and Owen Bowcott. Gary McKinnon will not be extradited to US, Theresa May announces. *The Guardian*, 2017.

[91] UK Government. The UK Cyber Security Strategy, 2011. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/60961/uk-cyber-security-strategy-final.pdf.

[92] US Department of Justice. Department of justice announces new policy for charging cases under the computer fraud and abuse act, 2022. URL: https://www.justice.gov/opa/pr/department-justice-announces-new-policy-charging-cases-under-computer-fraud-and-abuse-act.

[93] Anh V Vu, Jack Hughes, Ildiko Pete, Ben Collier, Yi Ting Chua, Ilia Shumailov, and Alice Hutchings. Turning up the dial: the evolution of a cybercrime market through set-up, stable, and covid-19 eras. In *Proceedings of the ACM Internet Measurement Conference*, pages 551–566, 2020.

[94] David Wall. *Cybercrime: The transformation of crime in the information age*, volume 4. Polity, 2007.

[95] David S Wall. Towards a conceptualisation of cloud (cyber) crime. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, pages 529–538. Springer, 2017.

[96] Brad Wardman, Michael Weideman, Jakub Burgis, Nicole Harris, Blake Butler, and Nate Pratt. A practical analysis of the rise in mobile phishing. In Ali Dehghantanha, Mauro Conti, and Tooska Dargahi, editors, *Cyber Threat Intelligence*, pages 155–168. Springer International Publishing, 2018.

[97] Robert N. M. Watson, Simon W. Moore, Peter Sewell, and Peter G. Neumann. An introduction to CHERI. Technical Report UCAM-CL-TR-941, University of Cambridge Computer Laboratory, 2019.

[98] Marleen Weulen Kranenbarg, Thomas J Holt, and Jeroen van der Ham. Don't shoot the messenger! a criminological and computer science perspective on coordinated vulnerability disclosure. *Crime Science*, 7(1):1–9, 2018.

[99] Laurie Williams, Gary McGraw, and Sammy Migues. Engineering security vulnerability prevention, detection, and response. *IEEE Software*, 35(5):76–80, 2018.