Silicon Den: Cybercrime is Entrepreneurship

Ross Anderson, Universities of Cambridge and Edinburgh

Rainer Böhme, University of Innsbruck

Richard Clayton, University of Cambridge

Ben Collier, University of Edinburgh

Abstract

We present a fresh perspective on cybercrime: the viewpoint of entrepreneurship. We propose a framework that sets out what infrastructure enables a particular cybercrime to get started, what barriers to entry there may be, how the crime can be scaled, what factors can inhibit scaling, why defenders can be ineffective and what eventual limits there may be to growth. We flesh this out via a series of case studies and draw out a number of common factors that cause cybercrimes to succeed, to scale and to stabilise. In addition to explaining the success (and failure) of particular cybercrimes, our approach may help predict the likely success of new types of cybercrime and the changing risks posed by evolution of the types we already know.

1 Introduction

Over the first two decades of the 21st century, online crime has grown from a niche phenomenon to dominate most measures of crime and victimisation – particularly, but not exclusively, in developed countries. The term 'cybercrime' now applies to a wide range of organised forms of technologically-facilitated crime and harm, which now exist in a complex ecosystem, yet are still only loosely understood as distinct social phenomena [54].

Within the large and expanding field of research and practice that has emerged around cybercrime and online harm, there are four broad approaches. The first is *technological*: engineers who design products and systems to resist attack collect taxonomies of harm and analyse them in terms of the 'cyber-attack chain', decomposing each particular attack into its component actions and then looking for one or more control points where links in the attack can be disrupted [4]. This kind of analysis may lead to a recommendation for better technical defences.

The second is *empirical*. It comes from the Internet measurement community and collects data on crime and abuse, giving us a picture of scale and concentration [23]. This may be used to set police priorities, but it crosses the boundary with our third perspective, the *economics of information security*, which has made the observation that the costs of cybercrime have remained largely constant for a decade, despite a huge technological shift from laptops to phones, from on-premise systems to the cloud, and from email to social media [5]. The problem is not therefore simply technological. Security economics also teaches that the failures of many large complex systems are due at least in part to misaligned incentives, where one principal pays for the defence but another bears the cost of failure [78]. This insight may lead to recommendations for improved governance.

The fourth approach is that of *criminology*, which takes sociological and psychological approaches to studying the people who commit crime; how they got involved, how they justify their actions to themselves and others, and why they eventually desist. Criminologists also study cybercrime and how societies react to it as social phenomena in their own right. Much of their work has involved applying existing 'classic' criminological theories to cybercrime, with mixed success; different theories give better accounts depending on whether the perpetrator's core skills are technical or social, and on the stage that they have reached in their criminal career [49].

Ultimately, however, a compelling characterisation of 'cybercrime' remains elusive. Ideally we would like to encompass everything from Internet-facilitated scams to the development and deployment of complex malware, and understand why some forms of cybercrime succeed while others fade away. So far we have perspectives obtained by considering attack methods, defence options, defence priorities, the systemic organisation of defences, and the social reasons why particular individuals may choose or be pushed towards a criminal career.

However an important perspective is missing. Many of those involved in cybercrime are entrepreneurs, looking to do something different: whether to invent a new scam, to take a legacy crime online, to improve an existing cybercrime, or at least to perform an illegal activity that they believe to be profitable in the hands of others. It is these entrepreneurs who turn ideas into businesses – who take social and technical vulnerabilities in human systems and manage to create stable formations of practice around them in ways which mirror 'disruptive' innovation in other sectors.

Not all of those involved in cybercrime are truly entrepreneurs; just as in the regular economy, successful entrepreneurs may hire others to do the boring work [24]. And not even all cybercrime innovators are entrepreneurs. For years, malware was illegal but essentially a sport, as young programmers wrote it for hacker cred – to show off their skills and gain status. It was only when entrepreneurs figured out how to use malware to take over machines to send spam that malware became a business. Then, once professional malware was available, people found new applications, from a variety of attacks on banking systems [107] to ransomware [92].

Of course, many cybercrime business owners are entrepreneurs in the fairly weak sense that they are copying a well-known modus operandi, rather than pioneering a new business model. Again, this has parallels in the world of regular business; most self-employed people and small company owners follow a standard pattern. They may copy other businesses that they see, or fit into a more formal structure such as a guild, a profession or a franchising system.

Scholars of entrepreneurship and industrial organisation already study what's involved in starting a business, whether it's a high-risk tech startup or a lower-risk venture such as a restaurant or shop. In the same way, scholars of cybercrime are beginning to understand why and how some types of activity scale. The objective of this paper is to organise this knowledge, to look for patterns, and to see what we can learn. At this stage our objective is understanding, although in the medium term we would like to be able to predict whether some new crime has the potential to scale up into a real problem – and if so, what technical, social or other controls might be most likely to hinder its growth.

2 A new approach

The main contribution of this paper is to look at cybercrime through the lens of entrepreneurship. Although many have observed that cybercrimes are businesses operating in a market, there is much more to it than that. There is a rich literature on entrepreneurship, taking risks to set up a business and make a profit, but it has hardly been applied to the study of cybercrime. Some topics are of little relevance to a criminal enterprise – location decisions, the mechanics of venture funding, intellectual property strategies, how to exit via an IPO or trade sale – but much of the discussion about what makes an entrepreneur successful is of direct relevance to cybercrime and cybercriminals [2, 15, 29, 32, 40, 47].

There is a much smaller literature on the entrepreneurial aspects of traditional crime. In addition to the archetypes of the middle-class entrepreneur and the working-class boy or peasant who made good, there have long been criminalised entrepreneurs such as in sex work or the sale of illegal drugs [94, 44, 18]. Criminal entrepreneurship is mixed in with some businesses; food supply crime can involve rogue farmers selling stock to unlicensed slaughtermen [75] and wildlife crime such as illegal, unreported fishing, which is significantly correlated with poor governance [3]. However the only two works on online criminal entrepreneurship of which we are aware are Kraemer-Mbula, Tang and Rush's study of the value chain in credit-card fraud [60]; and Flamand and Décary-Hétou's study of online drug dealing [39]. There appears to have been no work on the technological crime aspects.

The interaction between entrepreneurship and the state might be one starting point for an analysis of online criminal enterprises. Baumol laid the foundations with a study of why many more businesses are started in some societies (Europe, from medieval times onwards, and the USA) than in others (such as ancient Rome and imperial China) [11]. Lerner then studied the difficulties that governments in both Europe and America encountered in promoting technology clusters [64], while Janeway analysed the historical interaction between government and tech entrepreneurship in the Bay Area and elsewhere [55], documenting how government funds most of the science on which such ventures are based, and how disruptive innovators often bump up against laws made to protect legacy industries and the people who work in them.

But while firms like Uber can raise billions of dollars in venture capital to challenge taxi regulation in hundreds of cities, the businesses of interest to us are those whose goals and methods place them beyond the contestable grey area at the edge of the law. This generally constrains them to start on small budgets and grow via retained earnings.

The peer-reviewed subset of the entrepreneurship literature is somewhat chaotic, with multiple incompatible definitions and models [80]. Our approach has therefore been pragmatic.[1] Our empirical starting point is the case studies to be found in twenty years of papers that measure cybercrime, which have appeared at conferences such as USENIX Security, APWG eCrime, IMC and of course WEIS. We have chosen our examples to be illustrative rather than exhaustive.

Cybercrime has been more or less stable for the last ten years [5] so we have taken the view that rather than looking at technical or social factors separately – or proposing yet another taxonomy on the basis of skill, tech, and other such factors – we should be looking at the stable formations which have now emerged and looking for common factors between them.

The framework we have developed sets out what is necessary for a particular cybercrime to succeed, in the sense of becoming relatively stable at scale; in other words, a volume crime. The six elements of this framework are synthesised from our own observations and analysis, the cybercrime research literature, and concepts from the entrepreneurship literature which explain why particular tech businesses scale (as can be found in, for example, Suzuki et al. [105])

Applying this framework to a selection of current and historical cybercrimes allows us to identify common factors that appear to be important for a crime type to get started, to scale up (whether using infrastructure or by recruiting more participants), and to escape effective action by law enforcement. Further factors include whether a scam's growth is sustainable, or whether there is some natural limit; in some cases, people no longer fall for a scam once everybody has heard of it, while in others (such as underground drug markets) network effects create a 'success disaster' as they drive all the traffic to a single enterprise which becomes so prominent that it gets taken down by law enforcement – or is unable to scale up its services quickly enough to meet demand and collapses.

3 Framework

Our framework has six steps, illustrated in Figure 1:

1. preconditions that enable a particular type of crime;
2. barriers to entry that entrepreneurs must somehow overcome;

[1] Three of the authors have started companies; two have sold them; and one has been an angel investor.

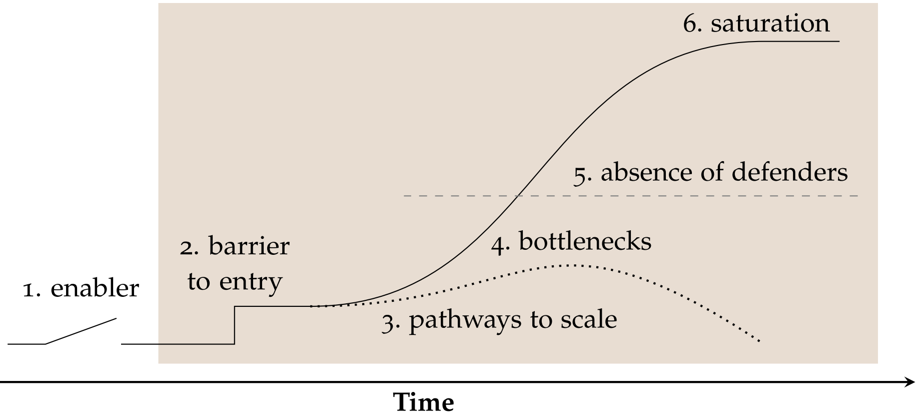

Figure 1: Trajectory of a new crime type

3. pathways involved in scaling up the crime to something of consequence, whether to the perpetrators, to society or both;
4. bottlenecks that can inhibit scaling, and potentially act as control points;
5. an absence of motivated defenders, as may happen if incentives are misaligned;
6. and finally, saturation mechanisms that set natural limits to growth.

As an illustrative example, consider '419' advance fee fraud, in which marks are offered a share of an illicit fortune in return for helping to launder the funds, and are then faced with demands for 'bank fees' or 'taxes'. This is a very old scam, known as the 'Spanish prisoner' fraud in the days when it was conducted by letter through the postal service.

The precondition for its becoming widespread (and a cybercrime) was that email enabled scammers to send millions of missives at near zero cost. The barrier to entry was low, namely access to a PC and to lists of email addresses; the pathway to scale was to recruit lots of young men to send spam, and later to send it automatically; the bottleneck was the development of spam filters; there was an absence of motivated defenders among the police in both sending and receiving countries, as the sums involved were mostly below the threshold for international cooperation; but eventually the crime saturated once everyone had heard of it.

Just as the series of technical steps required for a computer crime can be considered a 'kill chain' which the defender can try to interrupt, so also our framework gives a 'scam chain' of the business conditions and processes necessary for a new crime type to scale up to the point where it becomes a nuisance. We will now examine the steps of the scam chain in detail, and support it with relevant case studies.

3.1 Preconditions and enablers

The first factors we consider are the preconditions and enablers which create opportunities for criminal businesses to emerge. These are relatively stable resources, both social and technical, that overcome barriers to entry and support other steps of the crime pathway, particularly scaling. Regular businesses generally rely on four such supportive enablers: technical infrastructure, markets, finance, and governance.

Infrastructure consists of stable arrangements of technologies and supportive practices on which higher layers of services and technologies can be built. As infrastructures develop over time – such as the water and electricity grids, railways, roads, and the telephone network – they allow particular services and functions to be 'taken for granted' by society, permitting further services to be layered above them [101, 102]. Digital infrastructures, initially built on top of existing computers and telephone communications, have been revolutionary as they enable higher layers of infrastructure to be built out of software and deployed at very low marginal cost. Two examples are the Internet, whose TCP/IP protocols enable higher layers to disregard the mechanics of how two devices are connected, and the World-Wide Web, with its universal protocols for handling and displaying a range of content types. These massively reduced the barriers to innovation and enabled the mass deployment of further and more complex infrastructures from the 1990s [28].

General-purpose infrastructures also allow criminal practices to function and scale. In the digital domain, these include communications services, such as email and VOIP; payment services such as PayPal and online banking; and social networking services like Facebook, Twitter, and Discord. As with roads and railways before them, these enable crime and harm in various ways, leading to a game of cat-and-mouse between the police, the owners of the platforms, and those who abuse, exploit or subvert them.

Other infrastructures are more contested in that they make many crimes easier, but are still tolerated because they also have substantial non-infringing uses. Tor and Bitcoin are the most obvious examples. A recurring political debate is about whether end-to-end encrypted chat should be forced to allow law-enforcement access.

Finally, some illicit infrastructures are primarily designed to support crime. The most obvious examples are botnets and underground crime forums; there are also unlawful services from booters to unlicensed cryptocurrency markets. Some legitimate services can also become deviant, as when offshore call centres turn to selling unregulated investments or to outright phone scams. Additionally, there is 'scrap infrastructure' – legitimate infrastructure that is so defunct, insecure, and poorly maintained that it can be easily subverted, functioning as a de facto criminal infrastructure. An example of this is the large numbers of badly-secured Internet-connected devices around the world, which can easily be turned into denial-of-service reflectors [24].

Besides technical infrastructure, we need to consider markets – institutions that match the supply of labour, victims, and customers with the demands of criminal enterprises.

Given technical infrastructure and markets, we have some basic preconditions for an underground economy. The regular tech economy also has finance, from VC funding for startups through ways of selling your company to cash out, but the illicit economy generally lacks these. It also lacks governance, such as courts to enforce contracts and resolve disputes about liability. In the world of legacy crime, we see mafias emerging as alternative dispute resolution mechanisms, but cybercriminals must try to build governance mechanisms themselves (e.g., through escrow and reputation systems) or get by without them [73]. This is one of the reasons why cybercrime enterprises generally seem harder to scale.

We do not rush to make value judgements about the boundary between illicit and regular business. Some people consider a lot of legal entrepreneurship to be morally reprehensible (such as payday loans) while others campaign to legalise illicit entrepreneurship (such as the production and sale of marijuana). There are genuinely difficult cases in the middle for people of any political persuasion, such as the regulation of cryptocurrency services. For our purposes, the important distinction is that illicit entrepreneurs do not have access to some of the supportive services and infrastructures of the mainstream economy. A central point we want to make about cybercrime is that many of these enterprises are effectively trying to run an innovative tech start-up or digital business with the financial structure of an ice cream shop!

3.2 Barriers to entry

The preconditions and enablers we just discussed are powerful factors in enabling cybercrime, as they overcome a variety of pre-existing barriers. One historic barrier was connectivity – the physical distance between illicit enterprises and the supply chains, customer markets and victims which they need to access. This was knocked away when then Internet allowed people to communicate worldwide at near-zero marginal cost. As Naughton puts it [82], the Internet has many of the aspects of a failed state – and we are all living just over the border from it.

Another barrier is skill. Technical hacking skills are highly transferable, as are sales and marketing skills. People who acquire social-engineering experience at one type of crime – or even through working in a legitimate sales office – can transfer this to other types of crime, from fake support scams to spear-phishing. Marketing skills can help in the design of websites and phishing lures. Skill barriers can be overcome not just by training but by hacking tools and crime scripts. Indeed there are now 'howto' guides for a number of crime types that tell a newbie exactly what they need to buy and to do – just as you can franchise a coffee outlet.

A related factor is contextual domain knowledge – knowing, for example, that most police forces won't prosecute criminals whose victims are all overseas, and what sort of offences they will actually investigate. This is crucial to risk-management decisions. As well as mere business risk, the criminal entrepreneur faces the risk of jail. Diving deeper, there are normative barriers – committing offences may involve violating social norms against harming others. So participants may need to internalise the norms of a particular criminal subculture and the ways in which potential harm to others is rationalised [106, 16]. (Mainstream business also has culturally embedded ways to rationalise harm to customers, to society and to the environment.)

This is not only a matter of norms around harm. The entrepreneurial mindset also includes skills and norms, many of which have wide social approval and are also taught through online communities – another enabler that helps overcome a barrier to entry [49].

A less obvious example of a barrier arises from network effects in well-established markets. Once eBay got network effects going, it was more difficult for others to get into the auction business; sellers want to sell in the market with the most buyers, who in turn will go first to the market with the most sellers. In exactly the same way, when an underground market like Silk Road gets network effects going, it's hard to scale up a competitor. The pattern here is that when a dominant player like Silk Road gets taken down by law enforcement, this creates an opportunity – and the next contender rapidly scales up to take its place.

The same pattern has been seen with forensic-resistant mobile phones. About once a year since 2015, the world's police forces close down the leading system, and within a few months another takes its place. We will discuss this further in section 4.11.

3.3 Pathways to scale

The ability to scale a business depends on several factors. First and foremost, economies of scale emerge if the fixed cost of production can be amortised over many units. In the context of cybercrime, the cost of finding a vulnerability and weaponizing it in malware, or alternatively of buying a ready-made exploit, is a nontrivial fixed cost, whereas the variable cost per attack can be as low as sending a single network packet (in the case of a wormable exploit). Variable costs do increase if some form of access control has to be circumvented; and they become significant if human intervention is required, such as social engineering to steal credentials.

Standardisation enables network effects between compatible systems, whereupon winner-takes-all dynamics favour a dominant platform [100]. The flip side includes homogeneous victim populations, which make it easy to scale attacks [9]. Diversity increases where humans are involved, whether as individuals or in organisations, leading to heterogeneous targets. This often means that only a small subset of all targets is going to fall for a particular scam. Consequently, scaling a crime may depend on one's ability to identify multiple targets in large populations with little effort [43].

The extreme end of target selection touches an important corner case where capable attackers go after high-value targets. This has striking parallels to an enterprise's strategic choice between cost leadership and premium product. Some crime types work both ways, as we shall explore for business email compromise (section 4.2) and ransomware (section 4.6). Many state-sponsored threats focus on well-defined targets, and try to reduce the risk of collateral damage, although their main reason for stealth is to keep exploits reusable rather than to reduce harm.

Scaling does not always mean increasing the number of users or victims. McCoy et al. [74] note that since tackling new market segments is expensive, repeat purchases are just as good for an illicit business as for a legitimate one. A criminal's reputation can be valuable, but is complex. Is he famous for keeping his word, or for taking revenge on his enemies? The former, at least, can be signalled via reputation systems.

Turning our attention to the defender, many defences could in theory scale to match the attackers' capabilities. Threat intelligence firms inform their customers not just about active malware but about the IP addresses of command-and-control servers, so their customers can identify and remove compromised machines at scale. And information sharing between defenders could scale up those defences that require a human in the loop [56, 103]. However, responders are usually fragmented within and across jurisdictions, and agile cybercriminals can stay ahead enough of the time.

3.4 Bottlenecks and off-switches

If one step of an attack chain fails to scale as well as the others, then a bottleneck emerges. This can help to limit the crime, at least until people innovate around it.

Clayton et al. [23] discuss concentration points found when cybercrime is carefully measured and understood. When bottlenecks are linked to concentrations, this is one place where defenders can make a difference. But concentrations arise for other reasons, and if they do not amount to a real bottleneck it may be easy to evade defenders' attempts to exploit them.

As an example, consider the takedown of McColo in 2008. This was a 'bulletproof' hosting company in California, which criminals used for child sexual abuse material, malware repositories, botnet command-and-control systems and spam senders. When the journalist Brian Krebs shamed the companies that connected McColo to the Internet into pulling the plug, all sorts of bad things disappeared from the Internet overnight [61]. But McColo's customers simply went elsewhere, and even the short-term impact was minimal; Clayton found that easy-to-block spam only decreased for a few days [22]. There was even less long-term impact, beyond a few spam gangs losing target lists they had not backed up.

A more important bottleneck can be liquidity. It is all very well stealing millions of credit card details, but if you plan to sell them on for others to exploit, then you rely on those others to buy them at scale and have the cash to pay for them. One problem is that cybercriminals do not have access to capital markets; when a juicy opportunity comes along, they cannot sell shares to a venture capital firm or take out loans in order to finance working capital. When 40 million sets of credit card details were stolen from Target in 2013, less than 3 million were sold on the underground markets – although over 21 million were reissued by the banks (and others would have expired) [14].

Herley and Florêncio suggested in 2009 that the price of credit cards on the IRC channels, which were important criminal marketplaces at the time, was so low because of rampant fraud (a "lemons' market"). An alternative explanation is a massive oversupply. Since prices have remained low, even though markets now have reputation and escrow mechanisms, oversupply seems to be a better explanation overall. The question is whether the market value of a credit card is exogenous (a function of how much value a crook could extract, and thus of technical protection) or set by marginal-cost economics (a function of whether there is a surplus of card data, or a surplus of cashout operators).

Other money-related bottlenecks arise if an attacker has money but cannot actually spend it. A key player in the Nigerian cybercrime community is someone with access to a first-world credit card. No matter the credit rating within Nigeria, a local credit card will not work for the purchase of domain names, website hosting or VPN services. So most Nigerian cybercriminals need to pay a middleman, which may create delays and slow down the scaling of attacks.

Other crime bottlenecks may be caused by the risk of detection. If a scam involves a phone company insider corruptly changing records for the scammer (as we will discuss later in section 4.10), they may not be prepared to do this very often for fear of being caught.

Off-switches that stop crime dead in its tracks are rare. Famously, the WannaCry malware was designed so that as soon as a particular domain had a working DNS record, it became harmless. When Hutchins followed standard procedure and registered a domain he found in the code (with a view to constructing a 'sinkhole') the threat ceased – except in parts of the world which had not understood the mechanism and were blocking the DNS traffic [72]. Another example was SQL Slammer in 2003. This was not cybercrime per se, but a worm that spread over UDP port 1434 – and blocking that traffic on transit networks was a significant mitigation. Similar blocking has been applied (though seldom discussed openly) to deal with denial-of-service attacks involving NTP MONLIST commands and memcached traffic.

Although off-switches are rare, self-limiting is common. The Mirai botnet was large because its innovative scanning algorithm let it outcompete earlier botnets for the same set of vulnerable machines. Scanning was an order of magnitude faster and so Mirai got there first when a device was rebooted – and then shut the door after itself [7]. However, once other botnet authors copied the Mirai code, the machine fleet was fragmented. Although new sets of vulnerable devices are identified regularly, the populations are usually small, and no single botnet approaches the size of Mirai.

Occasionally, competition for resources has been exploited to defeat the criminals. Smurf attacks[2] were obsoleted partly by changes to router software and partly by an incentive scheme to encourage patching. The criminals needed to scan the Internet to find the best 'smurf amplifiers', but a Norwegian researcher openly published a list of them. This seemed to make the criminals' task easier, but all the criminals then used the networks at the top of the list, maxing out their bandwidth and forcing them to apply the patches.

3.5 Absence of defenders

Scholars of security economics have noted that cybercrime often persists where defenders are unmotivated; when Alice guards a system but Bob pays the cost of failure, you can expect trouble [78]. In our introductory example of advance fee frauds, the perpetrator and the victim are in different countries and the sum is below the threshold for international police action. Few police forces care about volume crime committed by foreigners, and even fewer about volume crime committed against foreigners. The responsibility for defending against cybercrime also generally involves 'multi-agency partnerships' including private security firms, volunteer groups and private infrastructure providers [67]. US law-enforcement agencies spend more on fighting cybercrime than the next ten countries put together; most other governments free-ride [5].

Local forces generally lack the technical skills and international links to tackle volume online crime, while centralised agencies focus on national-security and organised-crime threats, considering the vast ocean of smaller-scale harm to be out of scope [122].

The private sector has mixed incentives. The tech majors spend about the same on fighting cybercrime as the US law-enforcement agencies, but care only about particular offences that damage their reputation or cost them money directly. Smaller entities and the general public often lack the capacity to protect themselves. Although industry and volunteer groups have had some successes in working together – such as in reporting suspect IP addresses to blacklists – private-sector efforts tend to reflect the motivations, priorities and concerns of industry rather than end users and victims [56].

Beyond this, platforms and infrastructure providers have been reluctant to get involved. They argue both that they are neutral service providers and that they lack the capacity to respond to crime at scale. However there is increasing public and legal demand for them to take responsibility for the harms they facilitate. For Internet Service Providers, who are already fairly well-networked with the security services in most nations, this has been a less painful experience than for social media giants whose

[2] An early form of denial-of-service using broadcast ICMP echo request packets.

operations span many jurisdictions, and who have been loth to get involved in the political business of moderating content (though are now increasingly wading into these issues) [104].

3.6 Limits to scale and saturation

Saturation is reached when the space of vulnerable targets is exhausted, and can limit the growth of a criminal activity in both time and space. The main difference between a bottleneck and saturation is that the former kicks in orders of magnitude earlier. In fact, a bottleneck often provides, or at least suggests, an efficient off-switch.

Saturation does not need to cut in abruptly, as decreasing marginal returns may appear gradually as easy targets are taken first. Such effects can be rooted in technology or human factors. An example for the former is a long tail of rare system configurations, which gets increasingly harder to hit with automated exploits. Similarly, a long tail of attitudes, behavioural traits or norms may make it increasingly harder to trick people into falling for a scam. This can be as simple as language, where victims who speak the criminal's local languages are harvested first. Other large markets are then opened by hiring locals, or by automatic translation, but few will bother to localize a scam to small and culturally distant populations.

Other barriers to scale are cases of shrinking opportunity caused by technological advances (dial-up modems disappear, call-sell operations vanish with free VOIP calls) or increasing public awareness of a scam.

In principle, a criminal's strategic choice to stay under the radar of public attention or law enforcement efforts would qualify as a limit to scale. There are some clear examples of people trying to keep their activities under the radar, but this is not universal. CyberBunker, a bullet-proof hosting provider with a no-questions-asked policy, opted for visible advertising rather than stealth and got away with it for some years. Terrorist recruitment must by its nature reach out to potentially disaffected youth, and protest movements achieve impact by being annoying.

Where barriers to entry are low and defenders are absent, as with advance fee fraud, unlimited criminal competition can drive profits to zero, but at the same time create enough crime to raise public awareness that makes the crime harder to pull off.

3.7 Organisational forms

Once established, particular forms of cybercrime appear to settle into distinct organisational forms which help shape their scale, success, and resilience to enforcement. Much of the criminological literature has attempted to import ideas of organisation from classic forms of organised criminal activity – archetypal formations such as the mafia or the terrorist cell [41]. However, we support the view that these generally overplay the severity and level of organisation of cybercrime, while underplaying important structural elements of online organisation [17, 66, 70]. Wall [121] observes that online criminal formations tend not to operate like mafias, adopting more distributed, disorganised, businesslike forms. We agree, and argue that his analysis should be taken further in exploring how online businesses are organised and how organisational dynamics shape their growth.

In terms of illicit organisation online, we observe three main cybercrime formations which accord with standard models for tech businesses. First, there are mature markets, with a diverse range of small and medium providers (often connected with customers through a shared marketplace infrastructure) [10]. Secondly, there are forms of crime typified by small groups competing with one another for customers, resources, and victims with profits either extracted directly from victims or in secondary markets (such as the sale of stolen credit cards) [50]. Thirdly there are well-networked 'supply chains', in which distinct groups operate in a complementary ecosystem – hierarchies emerge here not in terms of formal structures, but in relationships between businesses, such as franchising, outsourcing, disintermediation of different complementary services, and reselling [24].

Cybercrimes often progress through these forms as markets and enablers develop over time, with entrepreneurs reacting to these changing conditions. For example, denial-of-service attacks began as a fairly artisanal technical pursuit, with individuals finding new ways to overwhelm systems and deploying them ad-hoc. This generated a market for tools which could be traded or sold. Following the formation of a political movement in the form of Anonymous and Lulzsec, this changed into a form of mass protest, then into small groups generating traffic at scale using technical means [97, 98]. Once there was a market for denial of service, more efficient providers emerged in the form of booter services. Initially, these were advertised on a common platform (Hack Forums), with a diverse and fairly flat market. Once they were thrown off, the market fractured, with progressive law enforcement action serving to centralise it around a single dominant provider.

This is a good example of how enforcement, technological innovation, business model innovation, and broader social change can all shift the forms of organisation associated with a particular cybercrime – often transforming its relationships with markets, barriers to entry, bottlenecks, and pathways to scale. This all places the focus on the entrepreneurs and how they cope with changing market conditions.

4 Case studies

In this section we present case studies of different types of cybercrime, each with a brief overview of the mechanics, the relevant factors of our framework, and an assessment as to how each crime type might be changing over time. We also try and show the differences between each example and the next – criminal activity is very far from being uniform.

4.1 Booters and Denial-of-Service attacks

At present the main use of denial-of-service attacks, that overwhelm Internet connected systems with high levels of traffic, is in providing online game players a way to cheat by incapacitating rivals [86]. They are occasionally used for more serious crimes such as attacking business rivals [84] or for extortion [108]. The firepower may come from orchestrating a botnet to direct traffic directly to a victim [7], or a reflected UDP attack may be used in which a server originates small request packets spoofed to appear to be coming from the victim. This results in very large response packets being sent from a large number of misconfigured servers [110].

These denial-of-service attacks are widely available to the unskilled through shop-front websites called booters – or euphemistically stressers, since they purport that their only use is for 'stress-testing' one's own system. Prices start low; at present, $10/month would be a typical price for an unlimited number of attacks, each of which is capable of swamping consumer connections [58]. Higher fees will purchase bigger attacks that will disrupt any business or educational institution, unless it has installed specialist equipment or moved its operations to a provider with specialist protection against attacks.

There are few barriers to entry for booter customers, although Karimi et al. showed that when PayPal stopped being available, and hence cryptocurrency payments were required, the number of customers fell significantly [59]. Nevertheless, the number of attacks continues to increase over time [1]. Running the booter infrastructure itself is a straightforward sysadmin task, albeit there tend to be less than a dozen booters of any size at any given time, plus perhaps 30 or 40 others who are struggling to make a profit. There is a regular pattern of law enforcement action, every year or two, against the market leader of the time, where its website is taken down and its operators prosecuted [25].

Booters are run by individuals, or small groups of two or three people, and there is evidence that burnout (boredom) may be a significant reason why some booters shut down [24]. There are natural limits to growth, as there's a finite number of online game players willing to cheat, and little scope for extortion, as the defences available to businesses are effective even if they are not cheap. Even high-profile successes such as the recent attacks on the New Zealand Stock Exchange [108] have tended to be successful only in terms of disruption, not in terms of being paid, and have not attracted many other criminals to join in.

4.2 BEC: Business Email Compromise

For a case study of a crime that is not yet in steady state but continues to grow in importance, we turn to Business Email Compromise (BEC). This is a blanket term for several different types of fraud, conducted over email, which socially engineer the victim into sending money to the criminal [53].

One version involves impersonating the CEO or another senior figure, requesting an urgent wire transfer of many thousands of dollars – on the pretext that this request had been overlooked when they were in the office. Competent fraudsters will compromise a company email account to access online calendars and understand email usage patterns. Then they wait until the CEO is known to be away from the office, and will include strict instructions not to disrupt the CEO's meeting by ringing their mobile phone.

A far less technically complex form of BEC involves just sending email impersonating the CEO to lower level employees asking for gift cards to be purchased 'as a surprise bonus for staff' and for the all-important serial numbers to be sent by email reply. Here there is generally little attempt at reconnaissance, but the email may be sent en masse to every possible employee – and if the criminal is lucky more than one will reply.

More sophisticated attacks use visibility into a compromised email system to identify high value invoices. Replacement invoices are forged with the criminal's bank details and sent to the victim along with a request to ignore the earlier version. Lookalike email addresses are used so that further conversation about the replacement invoice goes via the criminal in a 'man in the middle' attack on the payment process. In the shipping industry, these crimes have raked in millions – the value of entire cargoes [34].

There are few barriers to entry for the simplest forms of BEC, although skill is required to breach mail systems; a spear-phishing attack that collects credentials from a relevant person in IT, accounts or the C-suite may be a typical entry point. Successfully operating a 'man-in-the-middle' conversation between two parties to a business transaction requires domain knowledge about how that industry operates, but this is not very hard to acquire.

Some ad-hoc defences have been deployed. Sellers of gift cards may explain the scam to would-be purchasers; some banks try to detect fraudulent transactions and delay them; security training may cover this type of fraud; accounts departments may check account changes carefully; and companies increasingly specify bank details in contracts, making them essentially immutable.

Reporting levels are high, because recovering funds is often possible if the fraud is identified quickly enough and that data shows that this crime is continuing to grow, year on year. Some of the growth is occurring because the crime is spreading to new forms of high value transaction – such as real-estate purchases [36].

4.3 Telecoms fraud

An example of a crime in decline is telecoms fraud. As Anderson et al. noted in their 2019 survey paper, headline figures from the Communications Fraud Control Association had dropped 23.2% from 2016 to 2017 and there is a further decrease in the latest report, which gives the 2019 figures [5]. A key reason for the decrease is that calls are now a lot cheaper so failing to pay for them involves less money! However, revenues are also reduced, and the percentage loss from fraud is now increasing. Not all telecoms fraud is cybercrime, but four of the fraud methods that are listed as "emerging"[3] are clearly cybercrime-related [19].

[3] "IP PBX Fraud, Abuse of network device, configuration weaknesses, and IOT Fraud"

Compromising Internet-connected telecoms equipment merely requires running off-the-shelf code to locate devices and brute-force weak passwords, but monetising such access is complex, often involving calls to premium-rate numbers in other countries. This crime does not seem to be well-documented, either in the academic literature or on underground forums.

In passing it should be noted that a related cybercrime, rogue-dialing fraud, has disappeared altogether. Malware would use dial-up modems to call premium rate numbers but changes to the revenue model – delaying settlement payments for several months to allow evidence of criminality to surface – pretty much killed it off, and in any case dial-up modems are now almost entirely a thing of the past [89].

4.4 Credit-card fraud

Anderson et al. also noted that although credit card fraud was growing, it was decreasing as a proportion of transaction volume and they concluded that this should be seen as a success [5]. This trend is confirmed by the most recent UK figures (for 2019) which show an increase in incidents but a decrease in value per incident – suggesting that fraud was detected better but more cards were falling into criminals' hands and it was easier to scale up fraudulent activity [113].

Driving the increase in incidents is the trend away from individual compromises of credit card details to bulk acquisition through compromise of either payment terminal systems or merchant websites. There are some significant barriers to entry here in the level of expertise needed to carry out such attacks.

Once credit-card details are compromised, the limiting factor on the amount of fraud that follows will be partly the ability of the bank fraud teams to detect unauthorised use, and partly the liquidity in the marketplaces where the cards are sold on to those who will actually do the fraud.

These underground marketplaces have been widely studied, often with a view to understanding the mechanisms, such as feedback, which allow criminals to place some trust in the marketplace and each other [46]. Another strand of work looks at the opportunities for disruption and intervention in these markets [51], although there is little if any research into which interventions might actually be effective.

4.5 Phishing

Phishing, which we define as social engineering attacks over email (or latterly SMS) that induce disclosure of credentials, has a long history. The days when the 'rock-phish' gang sent billions of emails a day [76] have long passed – sophisticated spam filtering systems at mailbox providers have made it very difficult to deliver malicious email at anything like this scale. Individual attacks tend to be at much lower volume and this, along with improved detection techniques has led to a substantial rise in the number of unique phishing sites identified each month [8].

Little technical skill is needed to attempt a phishing attack, but many attacks are poorly executed and have little impact. Nevertheless, if a criminal is able to develop a realistic-looking phishing page (or has the capital to purchase one from a specialist developer) then they can be very successful – Oest et al. recently reported an overall 7.42% success rate against a major financial institution [87].

Useful measurements of phishing have long been recognised as extremely difficult [77]. However, it is much easier to see that the targets have changed radically. Fifteen years ago almost all phishing was directed against banks, but now the main targets are generally PayPal, Amazon, NetFlix and the largest email providers. This has been brought about by the banks' use of two-factor authentication, plus sophisticated fingerprinting of machines and browsers. These changes are expensive, and cause some inconvenience to legitimate customers, so adoption is patchy.

Although the European Banking Authority has mandated two-factor authentication, many banks use phones as second factors to save money, and where the customer uses a banking app on the same phone, banks are permitted to consider a second app as a second factor rather than demanding (or issuing) a second device. Nevertheless, measures that go beyond just accepting a password as proof of identity remain an option for any company or industry that can no longer tolerate being attacked by phishers.

4.6 Ransomware

Ransomware might be the most compelling case of an enabler; the most difficult thing about kidnapping is not getting caught as you pick up the ransom. People tried gift cards for payment in the mid-2000s, but this could not scale to million-dollar ransoms. What made ransomware viable was cryptocurrency.

The basic mechanism of ransomware is that a victim's data is encrypted or downloaded and a ransom must be paid to recover it, but there has been a marked change in focus in the past few years. The highly successful CryptoLocker (2013–14) was randomly targeted and collected a few hundred dollars from each victim [69]. This resulted in a lot of copycat activity, but in the past few years there has been a major shift with criminals targeting organisations rather than individual victims, using far more sophisticated phishing attacks or machine compromises, and demanding tens of thousands, or even millions, of dollars in ransom [26]. Some argue that cyberinsurance has exacerbated this by making it more likely that victims will pay [33].

The criminals seem to have concentrated on particular sectors of the economy over time, first local government, then the health sector and more recently education. Since the victims often have backups, the criminals have taken to spending several days after they first gain a foothold installing full disk encryption, so as to ensure that multiple generations of backup are encrypted. They also frequently exfiltrate a copy of the data and then threaten to make this data public [26]. Releasing the data may be a regulatory issue or may significantly damage relationships with customers or suppliers. In one case, the criminals who exfiltrated more than 40 000 therapy records of psychiatric patients in Finland made contact with the individuals involved, offering to keep their records off the Internet if they paid €200 in Bitcoin [95]; the company involved later filed for bankruptcy [109].

Building ransomware, distributing it, accepting payments and then providing keys is a complex process. The CryptoLocker criminals found that they needed to spend a considerable amount of effort on customer support to assure victims that payment would actually mean that they got their data back [112]. The barriers to entry did not deter the unskilled – and a lot of the copycat systems have been deeply flawed, so decryption could be easy if one did not pay, or impossible even if one did.

By now, many attackers are highly skilled and law enforcement agencies have not been particularly effective at discouraging them. The way that the crime has evolved means that even companies that are well-prepared for a data loss are still vulnerable to extortion, so it is unclear what sort of 'saturation' will prevent this crime continuing to grow. Some high-profile attacks have not merely extracted seven-figure ransoms but affected supply chains: an attack on the Colonial pipeline interrupted fuel supplies to the Southeastern USA for seven days in May 2021, while another on the meat packing firm JBS disrupted some meat supplies in Australia, Canada and the US in early June. As a result, the US government has ordered law-enforcement agencies to treat ransomware as seriously as terrorism.

4.7 EWhoring

Pastrana et al. investigated the profits being made by EWhoring, which is the name given by one online community of criminals to the simulation of cybersexual encounters for financial gain [93]. The perpetrators impersonate young women, engage victims in social interactions on chat forums and then sell intimate photos and videos that they claim are of 'themselves' – but which have often been stolen from insecure online repositories [52].

The barriers to entry are low – and there are training documents and YouTube videos to explain effective techniques for the social engineering aspects. When packs of photos are needed these can be found for sale on underground forums – often by the makers of the how-to material.

The crime is narrowly defined by the training material. It might seem a small step to escalate the chat into persuading the victim to expose themselves online – and then blackmail would be possible. This is strongly discouraged not, perhaps, because it 'ups the ante' and might make law-enforcement action more likely – but because if a criminal operates a blackmail scam, then they will not need to purchase further packs of pictures.

Measurements of EWhoring activity suggest that it might be becoming slightly less common, but it's also possible to view it as having been in steady state for most of a decade. Although perpetrators and victims change over time, the crime itself is so defined and circumscribed by the photo pack suppliers that it cannot evolve. Perhaps the only real prospect for disruption is the advances being made in real-time video processing – you can already present yourself as a somewhat stylised cat [123] or add beard and glasses to your appearance. So we may see this crime displaced by suppliers of convincing overlays of young women, rather than mere photographs.

4.8 HYIPs – High Yield Investment Programs

High Yield Investment Programs are Ponzi schemes [79]. Their websites offer ludicrously high interest (1% or more per day), payment of which is financed (for a time) by new investors. There is a complex ecosystem of reputation websites listing newly established HYIP sites (on payment of a fee) and underpinning it all is a dominant supplier of turn-key websites (and indeed reputation websites as well) [83].

For several years the pitch made by the HYIP websites has been that investments are being made in bitcoin, as this cover story is more appropriate to the times than previous claims about investment in the oil business or foreign currency trading. There is a very low barrier to entry – you merely need enough capital to purchase an appropriate domain, website design and advertising on one or more reputation sites.

This is another crime which appears to be in steady state, for much the same reasons as we saw with EWhoring: little of no action by law enforcement, an ecosystem that supports a narrowly defined view of how the crime should operate, and a steady stream of new criminals and victims [93].

4.9 West African advance fee fraud & romance scams

We will discuss advance fee fraud and romance scams together, although they are very different types of crime. We do this because both are strongly associated with West Africa, and Nigeria in particular.

Advance fee fraud, often called '419' scams after the relevant section of the Nigerian criminal code, has moved on since the days where victims would find themselves contacted by people pretending to be princes (or dictator's wives) seeking to move funds out of their country. Most common these days are dying widows seeking to disburse their savings to charity, and – if the victim is sucked in – they will meet a wide range of bit part players, the hospital doctor, the catholic priest, the law firm that will draw up the paperwork, the barrister who will get a High Court certificate, the police officer who must look the other way and Interpol – who will want some anti-money laundering forms filled in. All of these must be paid, or bribed – all fees to be paid in advance of the victim laying their hands on the widow's money.

There appears to be a thriving economy of diverse specialists. Those who send the original emails sell on the hot prospects to someone else, who in turn will go to specialist suppliers of official-looking PDFs with the correct details – and to someone with a first-world credit card who can purchase the domain names and web hosting required for more elaborate scams. There is often a coordinator in charge of the whole scheme, directing who does what and when to cue the next actor onto the stage.

Romance scams, on the other hand, appear to be operated by small groups or individuals. They start with a lonely hearts advert, send an initial email, engage with victims over long periods of time, and occasionally make themselves very substantial sums of money [124].

Although these scams have been much studied from the victim's perspective, there is little known about the perpetrators – beyond what comes out in court cases when particularly successful practitioners are arrested. The studies occasionally done by security companies – in search of a blog post, a headline, or a talk at RSA – usually identify organised groups operating at some scale. This may be reporting bias, in that larger-than-life characters tend to be easier to identify and, when clues are followed up at random, the better-connected the gang the more likely that the clues lead to the same place and then a 'concentration' is identified that is perceived to be newsworthy. There are clearly some links to organised real-world gangs but it is hard to rate their importance relative to the large number of individual entrepreneurs.

4.10 SIM swapping

In SIM swapping, a malicious actor or group arranges for a provider of mobile telephone services to redirect calls and texts away from the victim to a device they control, by changing the SIM card associated with the target phone number. The swap is generally done by social-engineering a telco employee into believing that the new SIM is owned by the victim ("I lost my old phone, please will you help me transfer my number to the replacement"), or occasionally by bribing them to sidestep the due diligence required to establish that a swap request is legitimate [63].

As countries moved, at different times, to authenticate bank account access using mobile phones, SIM swapping started to occur; cases were reported in South Africa in 2007 [96], but not until 2015 in the UK [111]. This crime has become especially high-profile in the past few years, once it was realised that the technique could be used to take over high-value cryptocurrency wallets [45]. It has also been used to take over social-media accounts, particularly ones with short or unusual handles – since these are valued in some communities, notoriously in the OG ("Original Gangster") underground forum [117].

Some of the threat actors associated with the OG community have started to use their social-engineering skills for other purposes, and took control of sysadmin credentials at Twitter, using this access to hijack VIP accounts and attempt a bitcoin scam [62]. This was too significant to be ignored and the perpetrators were arrested.

The key point with regard SIM swapping is that the telcos tune the operational security of SIM replacement to cover the risk to them of toll fraud. For example, Lee et al. found that security for post-paid accounts was better than for pre-paid accounts [63]. However these mechanisms are not enough to protect millions in cryptocurrency, or access to major corporate infrastructure. Other industries have therefore built key security mechanisms ('second factors') on sand. A second, and longer term, issue is that this cybercrime has bred a group of young people who have been honing their social engineering skills with little interference, until some activity became too blatant to be ignored. Their skills can be used for a wide range of other harmful activity.

4.11 Forensic-resistant phones

Some crimes connected with mobile phones do not have victims per se, but consist of the provision of criminal infrastructure. In March 2021, the Dutch and Belgian police took down Sky ECC, a forensic-resistant mobile phone network [85]. This was the latest in a series: in June 2020 the Dutch and French had taken down EncroChat [35], and that in turn followed Ennetcom [90], and a series of others. These products enable people to communicate using end-to-end encrypted messaging and VOIP, like WhatsApp or Signal, but are also hardened to stop the police who seize a phone from working out what other phones it has been in touch with. There is a huge market for these phones. When EncroChat was taken down it was initially claimed to have 70 000 users paying €3000 a year for the service, though this figure was cut to about 11 000 in later court papers.

The dynamic is that drug gangs and operators who support them (such as money launderers) buy these products; network effects cause them to converge on one particular brand; and this becomes such a prominent target for intelligence and law-enforcement agencies that they close it down. The usual law-enforcement playbook is to infect the phones with malware, then collect messages and traffic data for a surveillance period, followed by a technical takedown and mass arrests. A variant is to subvert the phone supply chain. In June 2021 the FBI and the Australian Federal Police announced that they had been running the Anom phone service for some time, after persuading its developer to cooperate in return for a reduced prison sentence. Anom had even picked up thousands of users displaced when the EncroChat and Sky systems were taken down [116].

This is good example of a market that fails to scale because of the police – or rather that when it scales, the police act. Other examples are botnets, and crimes requiring lots of foot soldiers, some of which will become informants. In fact, one of the roles of secure phones is to limit the damage that informants or arrested gang members can do.

5 Discussion

We discuss the broader insights gained from the application of this framework, starting with the criminal enterprise, then moving on to its environment and to criminal careers.

5.1 The cybercriminal entrepreneur

We have argued that many cybercrimes are best understood as forms of tech entrepreneurship, driven by the economic, social, technical, and organisational motivations of a tech business. This new perspective provides a powerful way to look at some of their key qualities. It allows us to make sense of why some offences become stable 'volume' crimes and why some remain niche pursuits (or collapse after changes in market conditions). By working through a set of examples, we have explored these entrepreneurial dynamics in a range of 'classic' cybercrimes.

One of the striking features is that almost all crime groups are small, consisting of a few buddies. It is well-known in the regular tech ecosystem that the costs of software development and maintenance increase rapidly once there are more people than 'can be fed from two pizzas', as Amazon founder Jeff Bezos memorably put it. Agile development assumes a team who can hang out in one room, or now perhaps on one video call. Multiple teams working on a single project necessitate interaction by multiple managers and everything slows down. In a criminal organisation the dynamics of trust, from the risk of informants to the risk of a group splitting into two competing organisations, make larger teams even less attractive. In reality, many cybercrime groups form in much the same way as rock bands – people link up with their siblings, people from school, people they admire, or people who answer a wanted ad (whether in a music magazine or on a cybercrime forum). Only in a very small number of cases does an industry manager come in and assemble a team to be the 'next big thing'.

So how can criminal groups achieve scale? A common historical approach was the extended family, as seen from the organised crime groups of southern Italy to the thuggee of nineteenth-century India [119]. There are at least two cybercrime examples: Romanian villages that have specialised in ATM crime, and Chinese villages involved in gaming scams [125]. This is one way of adding people while managing the risk of training up competitors.

Another approach is business-process outsourcing. Call-centre staff in India have been used for some years in fake AV and 'Microsoft support' scams, although the Indian police are now starting to make regular arrests of such operators [88]. Individual support staff can also be hired as contractors, particularly to do the boring sysadmin work of maintaining criminal infrastructure [24].

Yet another approach is franchising, which has precursors in legacy crime; see for example Levitt and Venkatesh on drug gangs [68]. There, franchises are for geographical territory that franchisees defend by force; the Internet largely lacks defensible space, so franchises must be enforced using technical mechanisms. In the EWhoring ecosystem we described in section 4.7, the franchising is around packs of pictures and crime scripts; with the high-yield investment programs we described in section 4.8, it's about the reputation system and the rest of the support infrastructure. There is also crimeware-as-a-service where authors of banking Trojans make them available to other gangs in return for a cut of the proceeds, while embedding in their software some mechanisms to monitor how many bank customers get victimised [114].

But can scaling be purely technical? One big difference with tech startups is that a small team that hits a sweet spot with a popular service can scale it very quickly using rented infrastructure. An early example was YouTube, created by three ex-PayPal staff in 2005, and sold to Google the following year for $1.65bn. YouTube was followed by WhatsApp, Instagram, and others. Another example was PlentyOfFish.com which offered free online dating, paid for by Google ads, and hit $30,000 a day in ad revenue by the time its founder, Markus Frind, hired his first employee in 2007 [71].

We do indeed see some cybercriminals scaling up businesses that make millions in annual profits from many small transactions, including rental scammers [118], operators of illegal payment systems [115], and cryptocurrency-based investment scams [6]. But most of the crime types we describe in section 4 above do not scale this way. Two of them – ransomware and business email compromise – make millions for their operators via a small number of high-value frauds. Only one of them (forensic-resistant phones) involves engineering at physical scale. In that case, the operator buys thousands of Android phones, maybe modifies the data port, installs software, packs them into their own brand of packaging and sells them by mail order or through existing criminal networks. A similar historical example was the sale of Viagra by mail order while it was still in patent. Those entrepreneurs made several million a year by buying pills from countries that didn't respect Pfizer's patent and mailing them to countries that, in theory, did. (The bottleneck turned out to be card payments, and the business petered out when the Viagra patent expired.)

Note that a lot of our crime types become generic over time as people innovate, adapt well-worn scripts, and mine out the space around a particular form of crime – although fifteen years ago we would have headed the section "Nigerian Princes", now it's generic "Advance Fee Fraud". Selling illegal goods by mail order is another generic type, and investment fraud is a third.

5.2 Supportive infrastructure

Many forms of cybercrime have changed and evolved over time. When a cybercrime exploit is initially discovered, it often starts out as a niche pursuit of individual artisans. But as more people become involved, those with human skills go after bigger prey (as with ransomware and BEC) while those with technical skills go after small prey at greater scale (as with booters, phishing for bank credentials and some investment scams).

When scaling up, an operator may abuse legitimate infrastructure, such as commercial hosting companies, or rent illicit infrastructure such as botnets. The former leads to problems when the providers eventually notice and try to kick you out, so you need to keep registering new domains, opening accounts at new hosting providers, and so on, and so on. Illicit infrastructure needs even more maintenance, and often requires a higher level of technical skill. The two are not perfect substitutes, and a number of crime types require both. There is a grey area between them of legitimate infrastructure that is poorly secured and managed, such as cryptocurrency exchanges that pretend to keep proper records but don't really. As in legacy crime, the operator needs to have enough situational awareness to know what they can reasonably hope to get away with, and what may attract a capable response.

A lot of the business of innovation involves taking well-tested existing forms of crime online, then scaling them up using supportive infrastructure to automate work, to create a service market, or to establish governance and trust functions – such as the reputation and escrow services found in many underground marketplaces.

These marketplaces have been around since at least 2004; they allow cybercriminals to specialise and become good at their jobs. Crooks no longer have to run a vertically-integrated business but can focus on writing malware, on harvesting credit cards or online banking credentials, on cashout, on denial-of-service, on sending spam and many other services. There is a range of such markets, some open and some invitation-only, some on the open web and others operated as Tor hidden services.

A potential future concern is unstoppable computer programs, so-called “smart contracts” [57]. At present, their main identifiable criminal use consists of automated cryptocurrency exchanges that are designed to keep no records, thus making it harder for the police to trace drug money that is switched from bitcoin to ethereum to litecoin (say). In theory, financial regulators should close down such operations but FATF took until March 2021 to rule that their operators need AML/KYC controls [20]. In practice, there may also be ways of manipulating automated exchanges, leading to new types of technical fraud [30]. If smart contracts become more dependable and widespread, they could provide automatic contract enforcement, thereby levelling out the big difference between legitimate and criminal firms: that only the former can get courts to enforce their contracts.[4] This could enlarge the criminal firm’s space of feasible contracts, thereby enabling new types of crime.

5.3 Cybercriminal careers and capacity transfer

Looking at one crime type at a time through our new lens is instructive, but it may hide how they are connected. The common enablers are not limited to technical infrastructure, but include spillover via transferable skills. Just as there are serial entrepreneurs who have participated in multiple tech startups, those working in cybercrime make careers, too. A growing number of court cases involve second or subsequent offenders, including the TalkTalk hacker Eliot Gunton; the JP Morgan hackers Gery Shalon, Joshua Samuel Aaron and Ziv Orenstein; the Liberty Reserve operators Arthur Budovsky and Vladimir Kats; and the TJ Maxx hacker Albert Gonzalez.

There are both technical and psychological aspects to this. The former may include familiarity with hacking tools, underground markets, criminal infrastructure and the operation of payment networks. The latter may include skill at social engineering, sales and marketing generally; and there’s also whether people are psychologically prepared to be entrepreneurs, to break the law, or both. Experience of self-employment is widespread but not universal, particularly in developed countries; many young people have no experience of working other then for a salary. Experience of petty online crime may be more widespread but we are aware of no real data. Role models matter; underground cybercrime forums can provide bad role models which amount to a pathway to crime. A youngster may start off as a gamer, then cheat at games, then buy game cheats, then trade game cheats, then move to operating a booter service, and end up working with malware [49].

The experience of running a business and the experience of committing online crimes are reflected in the norms of various subcommunities – both deviant norms around willingness to ignore or even justify the harmful consequences of one’s actions, and entrepreneurial norms around succeeding by making money. These can transfer between legitimate and illegitimate enterprises. Indeed there is a large grey area in between, encompassing such activities as spam, malvertising, vulnerability brokerage, deceptive consumer dialogues, search engine optimisation, and covert surveillance.

6 Previous approaches

Economic theory has been applied to crime since the eighteenth century with early classical theories of crime and justice [12], and came to further prominence in the 1960s thanks to seminal work by Becker [13] and subsequent developments of his Chicago-school approach [37, 38]. Their tradition analyses deviant behaviour in terms of individual rational actors making calculative decisions, with crime a product of an internal calculus of opportunity, cost, benefit and risk.

[4] Real-world mafias do provide enforcement mechanisms [38].

Widening the scope from individuals to groups, Schelling's definition of organised crime as a form of governance structure – rather than a mere business – has shaped decades of work [99]. But its applicability to profit-oriented cybercrime has been challenged [70]. If the loose relations between cybercriminals resemble a bazaar more than a mafia, then a good way of analysing a cybercriminal firm might be transaction cost economics [31]. Regardless of how an organisation is managed, a precondition is to understand from which pool it recruits. The part of development economics concerned with the allocation of talent thus offers a link to entrepreneurship [11].

One of the well-established descendants of the Chicago-school approach within criminology is Routine Activities Theory, which draws it out into ecologies of criminal opportunity [37, 91]. Crime is held to emerge where three features coincide – a motivated offender, a suitable target, and the lack of a capable guardian. It follows that large-scale social changes – such as the move of women into the labour market, the rise of cheap, lightweight consumer goods, and the retreat of the American middle class into the suburbs – create or destroy opportunities for crime.

The same analysis has been applied to make sense of the opportunities for crime and harm posed by Internet technologies [65], and more recently to understand how online crime is changing in the COVID-19 pandemic [48]. Two big changes are the multiplication of opportunities for crime brought by the Internet connecting potential offenders to potential victims all over the world, and the fragmentation of the jurisdictional powers the police use to trace offenders and secure arrests, reducing the presence of capable guardians [120].

Routine Activities Theory has been used to develop strategies for crime reduction by modifying the ecology of criminal opportunity. In its purest form, the Situational Crime Prevention approach aims to mitigate opportunities for harm by modifying the situations where criminal opportunities appear – often through target hardening measures such as security cameras in the built environment [21]. Such changes attempt to raise crime difficulty and risk while reducing potential rewards.

In accordance with modern approaches to crime and justice, these security measures are generally not provided by the police, but are sold in the market and perhaps forced on private firms and households via insurance mandates. For cybercrime, too, this approach has tended to place on potential victims much of the responsibility for security, such as password management and software updates [27]. Ultimately, these theories rely on an economic analysis where individual victims and offenders make rational decisions in their own interests.

Some criminological writing has already described forms of crime as 'deviant entrepreneurship' and noted that cybercriminals in particular operate businesses in markets [42, 44, 81]. Musotto and Wall have begun to discuss the business-like aspects of how cybercrime is organised, but focus on empirical assessment of the characteristics of group organisation and customer markets [81]. They contribute to a growing consensus within criminological scholarship that cybercrimes are poorly explained through existing frameworks for making sense of organised crime and mafias [70, 66].

However, so far nobody has tried to apply the many lessons learned in the regular tech economy and tech entrepreneurship. Yet these are our main source of insights into the effects of digital tools, technologies and infrastructures on how businesses scale. As we noted above, it is now much easier for small groups to build software services and infrastructures on top of the existing base of consumer devices, with the result that successful startups can achieve scale much more quickly than before.

In the criminal context, this means that where traditional illicit infrastructure such as drug running networks would require large numbers of workers and substantial investment in materials, digital infrastructure such as botnets and tools such as remote access Trojans (RATs) can be deployed far more easily and cheaply. If we want to understand why some forms of cybercrime succeed and scale, and why some falter and fail, we must acknowledge that the forms and determinants of success of the legitimate tech economy depend on similar properties. It therefore makes sense to use ideas about network effects from information economics; about bottlenecks and scaling from computer science; and about infrastructure from organisational theory.

7 Conclusions

Most of the acquisitive crime suffered by residents of developed countries is now online, yet our understanding of cybercrime remains poor and law enforcement is patchy at best. So far we have technical, economic and criminological approaches, and quite a lot of data. But we still do not have a coherent synthesis that accounts for technology, crime, enforcement, and organisation together.

In this paper, we have argued that entrepreneurship is one missing link. Cybercrime is more than just a business or a market – it is a particular type of business. It is much more akin to tech entrepreneurship than to our traditional models of organised crime, such as running a drug gang in South Chicago.

Cybercrime involves individuals or small groups of people creating online businesses. These businesses can then scale up using infrastructure provided by others, by franchising, or as a result of imitation by copycats. Along the way, there may be barriers to growth in the form of competing businesses; there may be bottlenecks, in the form of steps in the business process that require scarce resources such as people or access to payment networks. The business's growth may or may not be contested, depending on whether there are any motivated guardians; successful criminal businesses tend to be those that do not attract a rapid and vigorous response from the police, the banking industry or the tech majors. And finally, a criminal business may exhaust its market, at which point it may have to adapt or die.

This framework was derived from analysing a large number of existing cybercrime types, which we discussed above, as well as from our own experience of tech entrepreneurship. It may be of some value in a number of contexts.

First, it can better explain the dynamics of the cybercrime ecosystem. We know that overall the levels of crime and harm have been stable for about a decade, just like the regular business economy, yet individual offences come and go, as individual businesses do. Someone does a newsworthy scam; Brian Krebs publicises it; there's a rash of copycats; but once everyone has heard of it, it doesn't work any more. This is natural enough, but not something the established theories can easily explain. A business approach gives insight into the churn and why some crimes become well-established and start to scale.

Second, it can help us understand the effectiveness of different law enforcement actions. Recent work has suggested that targeting illicit infrastructure and people with specialist support skills (as done by the FBI) and targeting the demand market (as the NCA do with Cyber Choices) are two strategies that both have the potential to be extremely effective. Within our framework, this amounts to removing the infrastructure and market preconditions respectively.

Third, it can be used for prediction and triage. Next time the security industry says 'Everyone please panic now at this new cybercrime' we have a reasoned way to analyse and respond. How can it scale? Where are the bottlenecks? Are there hidden bottlenecks where criminals have been able to get away with abusing infrastructure or services? Where are the capable motivated opponents?

Fourth, it can be used for change tracking. Ransomware has become a different crime since about 2018 as capable operators learned to target firms rather then individuals, and it has now become a growing nuisance that may demand large-scale investment in system security. Our framework suggests that firms should invest in data-loss prevention tools as an early warning mechanism, and standardise a straightforward way for professionals to contact them when they detect the tell-tale signs of exfiltration. Another burst of growth might even persuade many companies that it is not such a great idea to let any and all of your staff access and download the records of 100m customers, which might do more for privacy than GDPR has managed so far.

Business email compromise is evolving slowly, but it is scaling up rapidly as competent attackers become more numerous and realise the wealth of targets. Things will not improve in the short term because this crime will only be tackled by revamping hard-to-change business processes – and modifying them will have a cost. You can embed bank account numbers into hard-to-alter contracts but you lose some of the flexibility of changing banks in search of lower fees.

Fifth and finally, our entrepreneurial analysis serves to organise the case studies into a coherent picture, while cautioning against simplistic approaches. A computer scientist might be tempted to see all cybercrime as a scaling war. Can the crooks scale up their operations in such a way that they don't attract a capable motivated response, while the cops optimise by looking for high-value targets? However the case studies show that it is a lot more complicated than that.

Some commentators are tempted to think of cybercriminals as being like a mafia. But cyberspace is very different because of the lack of hierarchy and territory. In the end, perhaps, we should just follow the money. Crimefighters traditionally understood this in terms of following crime proceeds to their ultimate destination; we advocate taking this further and analysing cybercrimes as business ventures. Cybercriminals generally don't have access to finance – and in the regular economy, you do not make big money by running a company, but by selling it, which involves external money, and usually at multiple stages. While a regular tech business can raise venture capital to get started, and cash out via a trade sale or IPO, these options are not open to crooks. In effect, they are trying to run tech startups, but with the financial arrangements of a street-corner ice-cream shop. This may be one of the answers to the old question of why, if the world depends on computers and all computers are insecure, the world has not collapsed [43].

References

[1] Ruba Abu-Salma. DoS attacks during lockdown: Worldwide data. COVID Briefing Paper 7, Cambridge Cybercrime Centre, 2020.

[2] Richard Adams, John Bessant, and Robert Phelps. Innovation management measurement: A review. *International Journal of Management Reviews*, 8(1):21–47, 2006.

[3] David Agnew, John Pearce, Ganapathiraju Pramod, Tom Peatman, Reg Watson, John Beddington, and Tony Pitcher. Estimating the worldwide extent of illegal fishing. *PlosOne*, 2009.

[4] Ross Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley, 3rd edition, 2020.

[5] Ross Anderson, Chris Barton, Rainer Bölme, 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, Ilia Shumailov, Mansoor Ahmed, and Alessandro Rietmann. Bitcoin redux. In *Workshop on the Economics of Information Security*, 2018.

[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] APWG. Phishing activity trends report: 4th quarter 2020. `https://docs.apwg.org/reports/apwg_trends_report_q4_2020.pdf`, 2021.

[9] Daniel G Arce. Malware and market share. *Journal of Cybersecurity*, 4(1), 2018.

[10] Angus Bancroft, Tim Squirrell, Andreas Zaunseder, and Irene Rafanell. Producing trust among illicit actors: A techno-social approach to an online illicit market. *Sociological Research Online*, 25(3):456–472, 2020.

[11] William J Baumol. Entrepreneurship: Productive, unproductive, and destructive. *Journal of Political Economy*, 98(5):893–921, 1990.

[12] Cesare Beccaria. On crimes and punishment, (trans. H Pallouci). *Indianapolis: Bobbs-Merrill*, 1764.

[13] Gary S Becker. Crime and punishment: an economic approach. *Journal of Political Economy*, 76(2):169–217, 1968.

[14] Natasha Bertrand. Here's what happened to your Target data that was hacked. `https://www.businessinsider.com/heres-what-happened-to-your-target-data-that-was-hacked-2014-10`, 2014.

[15] David L Birch. The job generation process. Technical report, MIT Program on Neighborhood and Regional Change, 1979.

[16] Shane Blackman. Subculture theory: An historical and contemporary assessment of the concept for understanding deviance. *Deviant Behavior*, 35(6):496–512, 2014.

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

[18] Sylvain Bureau and Jacqueline Fendt. Entrepreneurship in the informal economy: why it matters. *The International Journal of Entrepreneurship and Innovation*, 12(2):85–94, 2011.

[19] CFCA. Communications Fraud Control Association announces results of 2019 global telecom fraud survey, 2019.

[20] Chainalysis. FATF's proposed updated guidance for cryptocurrency regulation: Everything you need to know. `https://blog.chainalysis.com/reports/fatfs-updated-guidance-march-2021`, 2021.

[21] Ronald V Clarke. Situational crime prevention. *Crime and Justice*, 19:91–150, 1995.

[22] Richard Clayton. How much did shutting down McColo help? In *CEAS*, 2009.

[23] Richard Clayton, Tyler Moore, and Nicolas Christin. Concentrating correctly on cybercrime concentration. In *Workshop on the Economics of Information Security*, 2015.

[24] 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.

[25] 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.

[26] Coveware. Coveware quarterly ransomware report Feb 2021. `https://www.coveware.com/blog/ransomware-marketplace-report-q4-2020`, 2021.

[27] Cassandra Cross. 'Oh we can't actually do anything about that': The problematic nature of jurisdiction for online fraud victims. *Criminology & Criminal Justice*, 20(3):358–375, 2020.

[28] James Curran. Rethinking Internet history. In James Curran, Natalie Fenton, and Des Freedman, editors, *Misunderstanding the Internet*, pages 34–65. Routledge London, 2012.

[29] Michael S Dahl and Olav Sorenson. The embedded entrepreneur. *European Management Review*, 6:172–181, 2009.

[30] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *Symposium on Security and Privacy*, pages 910–927. IEEE, 2020.

[31] Andrew R Dick. When does organized crime pay? A transaction cost analysis. *International Review of Law and Economics*, 15(1):25–45, 1995.

[32] Matthew Dobbs and R T Hamilton. Small business growth: Recent evidence and new directions. *International Journal of Entrepreneurial Behavior & Research*, 13(5):296–322, 2007.

[33] Renee Dudley. The extortion economy: How insurance companies are fueling a rise in ransomware attacks. `https://www.propublica.org/article/the-extortion-economy-how-insurance-companies-are-fueling-a-rise-in-ransomware-attacks`, 2019.

[34] England and Wales High Court (Commercial Court). K v A [2019] EWHC 1118 (Comm), 2019.

[35] Europol. Dismantling of an encrypted network sends shockwaves through organised crime groupe across Europe. `https://www.europol.europa.eu/newsroom/news/dismantling-of-encrypted-network-sends-shockwaves-through-organised-crime-groups-across-europe`, 2020.

[36] Federal Bureau of Investigation. Public Service Announcement: Business e-mail compromise the 12 billion dollar scam. `https://www.ic3.gov/Media/Y2018/PSA180712`, 2018.

[37] Marcus Felson and Lawrence E Cohen. Human ecology and crime: A routine activity approach. *Human Ecology*, 8(4):389–406, 1980.

[38] Gianluca Fiorentini and Sam Peltzman. *The Economics of Organised Crime*. Cambridge University Press, 1996.

[39] Claudia Flamand and David Décary-Hétou. Chapter 3: The open and dark web. In *The Human Factor of Cybercrime*. Routledge, 2020.

[40] Pierre Giot and Armin Schwienbacher. IPOs, trade sales and liquidations: Modelling venture capital exits using survival analysis. *Journal of Banking and Finance*, 31(3):679–702, 2007.

[41] Misha Glenny. *Darkmarket: how hackers became the new mafia*. Random House, 2012.

[42] Petter Gottschalk. Entrepreneurship in organised crime. *International Journal of Entrepreneurship and Small Business*, 9(3):295–307, 2010.

[43] Cormac Herley. Security, cybercrime, and scale. *Communications of the ACM*, 57(9):64–71, 2014.

[44] Robert Francis Hesketh and Grace Robinson. Grafting: "the boyz" just doing business? Deviant entrepreneurship in street gangs. *Safer Communities*, 2019.

[45] David Hollerith. The rise of SIM swapping: how and why bitcoiners need to protect themselves. `https://bitcoinmagazine.com/culture/the-rise-of-sim-swapping-how-and-why-bitcoiners-need-to-protect-themselves`, 2021.

[46] Thomas J Holt, Olga Smirnova, and Alice Hutchings. Examining signals of trust in criminal markets online. *Journal of Cybersecurity*, 2(2):137–145, 2016.

[47] Gerke J Hoogstra and Jouke van Dijk. Explaining firm employment growth: Does location matter? *Small Business Economics*, 22:179–192, 2004.

[48] Shane Horgan, Ben Collier, Richard Jones, and Lynsay Shepherd. Re-territorialising the policing of cybercrime in the post-COVID-19 era: towards a new vision of local democratic cyber policing. *Journal of Criminal Psychology*, 2021.

[49] Alice Hutchings. Cybercrime trajectories: An integrated theory of initiation, maintenance, and desistance. In *Crime Online: Correlates, Causes, and Context*, pages 117–140. Durham: Carolina Academic Press, 2016.

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

[51] Alice Hutchings and Thomas J Holt. The online stolen data market: disruption and intervention approaches. *Global Crime*, 18(1):11–30, 2017.

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

[53] Internet Crime Complaint Center. Internet crime report 2020. `https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf`, 2021.

[54] Interpol, European Union, and Council of Europe. Guide for criminal justice statistics on cybercrime and electronic evidence, October 2020.

[55] William H Janeway. *Doing Capitalism in the Innovation Economy: Reconfiguring the Three-Player Game between Markets, Speculators and the State*. Cambridge, 2018.

[56] Mohammad Hanif Jhaveri, Orcun Cetin, Carlos Gañán, Tyler Moore, and Michel van Eeten. Abuse reporting and the fight against cybercrime. *ACM Computing Surveys*, 49(2), 2017.

[57] Ari Juels, Ahmed E Kosba, and Elaine Shi. The Ring of Gyges: Investigating the future of criminal smart contracts. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 283–295. ACM, 2016.

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

[59] Mohammad Karami, Youngsam Park, and Damon McCoy. Stress testing the booters: Understanding and undermining the business of DDoS services. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1033–1043, 2016.

[60] Erika Kraemer-Mbula, Puay Tang, and Howard Rush. The cybercrime ecosystem: Online innovation in the shadows? *Technological Forecasting and Social Change*, 80(3):541–555, 2013.

[61] Brian Krebs. Takedowns: The shuns and stuns that take the fight to the enemy. *McAfee Security Journal*, 2010.

[62] Brian Krebs. Who's behind Wednesday's epic Twitter hack? `https://krebsonsecurity.com/2020/07/whos-behind-wednesdays-epic-twitter-hack/`, 2020.

[63] Kevin Lee, Benjamin Kaiser, Jonathan Mayer, and Arvind Narayanan. An empirical study of wireless carrier authentication for SIM swaps. In Heather Richter Lipford and Sonia Chiasson, editors, *Symposium on Usable Privacy and Security*, pages 61–79. USENIX Association, 2020.

[64] Josh Lerner. *Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed–and What to Do About It*. Morgan Kaufmann, 2012.

[65] Eric Rutger Leukfeldt and Majid Yar. Applying routine activity theory to cybercrime: A theoretical and empirical analysis. *Deviant Behavior*, 37(3):263–280, 2016.

[66] Rutger Leukfeldt, Anita Lavorgna, and Edward R Kleemans. Organised cybercrime or cybercrime that is organised? an assessment of the conceptualisation of financial cybercrime as organised crime. *European Journal on Criminal Policy and Research*, 23(3):287–300, 2017.

[67] Michael Levi and Matthew Leighton Williams. Multi-agency partnerships in cybercrime reduction: Mapping the UK information assurance network cooperation space. *Information Management & Computer Security*, 2013.

[68] Steven D Levitt and Sudhir Alladi Venkatesh. An economic analysis of a drug-selling gang's finances. *The Quarterly Journal of Economics*, 115:755–789, 2000.

[69] Kevin Liao, Ziming Zhao, Adam Doupe, and Gail-Joon Ahn. Behind closed doors: measurement and analysis of CryptoLocker ransoms in Bitcoin. In *APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–13, 2016.

[70] Jonathan Lusthaus. How organised is organised cybercrime? *Global Crime*, 14(1):52–60, 2013.

[71] Richard Macmanus. Plentyoffish: 1-man company may be worth $1billion. `https://readwrite.com/2007/10/29/plentyoffish_one_billion/`, 2007.

[72] MalwareTech (Marcus Hutchins). Finding the kill switch to stop the spread of ransomware. `https://www.ncsc.gov.uk/blog-post/finding-kill-switch-stop-spread-ransomware-0`, 2017.

[73] Kimberley Masson and Angus Bancroft. 'Nice people doing shady things': Drugs and the morality of exchange in the darknet cryptomarkets. *International Journal of Drug Policy*, 58:78–84, 2018.

[74] Damon McCoy, Andreas Pitsillidis, Jordan Grant, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey Voelker, Stefan Savage, and Kirill Levchenko. Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs. *USENIX Security Symposium*, pages 1–16, 2012.

[75] Gerard McElwee, Robert Smith, and John Lever. Illegal activity in the UK halal (sheep) supply chain: Towards greater understanding. *Food Policy*, 69:166–175, 2017.

[76] Tyler Moore and Richard Clayton. Temporal correlations between spam and phishing websites. In *2nd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET 09)*. USENIX Association, 2009.

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

[78] Tyler Moore, Richard Clayton, and Ross Anderson. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, 2009.

[79] Tyler Moore, Jie Han, and Richard Clayton. The postmodern ponzi scheme: Empirical analysis of high-yield investment programs. In Angelos D. Keromytis, editor, *Financial Cryptography and Data Security*, pages 41–56. Springer Berlin Heidelberg, 2012.

[80] Peter W Moroz and Kevin Hindle. Entrepreneurship as a process: Toward harmonizing multiple perspectives. *Entrepreneurship theory and Practice*, 36(4):781–818, 2012.

[81] 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.

[82] John Naughton. Has the internet become a failed state? `https://www.theguardian.com/technology/2016/nov/27/has-internet-become-failed-state-crime-cyberspace`, 2016.

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

[84] Shaun Nichols. Brit hacker hired by Liberian telco to nobble rival now behind bars. `https://www.theregister.com/2019/01/14/liberian_hackerforhire_case/`, 2019.

[85] NL Times. Dutch cops take out encrypted chat service SkyECC; thirty arrests. `http://nltimes.nl/2021/03/09/dutch-cops-take-encrypted-chat-service-skyecc-thirty-arrests`, 2021.

[86] Arman Noroozian, Maciej Korczynski, Carlos Gañán, Daisuke Makita, Katsunari Yoshioka, and Michel van Eeten. Who gets the boot? analyzing victimization by DDoS-as-a-service. In *Proceedings*

of the International Symposium on Research in Attacks, Intrusions, and Defenses, RAID 2016, Lecture Notes in Computer Science. Springer, 2016.

[87] Adam Oest, Penghui Zhang, Brad Wardman, Eric Nunes, Jakub Burgis, Ali Zand, Kurt Thomas, Adam Doupé, and Gail-Joon Ahn. Sunrise to sunset: Analyzing the end-to-end life cycle and effectiveness of phishing attacks at scale. In *Proceedings of the USENIX Security Symposium (USENIX)*, 2020.

[88] The Times of India. Fake BPO busted at Delhi's Rajouri Garden, 2,300 duped of $1 million. `https://timesofindia.indiatimes.com/city/delhi/million-dollar-scam-over-2-2k-duped-in-garb-of-tech-support/articleshow/79089946.cms`, 2020.

[89] Ofcom. Ofcom approves amendment to ICSTIS Code of Practice. `https://www.ofcom.org.uk/about-ofcom/latest/media/media-releases/2005/ofcom-approves-amendment-to-icstis-code-of-practice`, 2005.

[90] Charlie Osborne. Dutch police close Ennetcom encrypted communications network. `https://www.zdnet.com/article/dutch-police-arrest-owner-of-ennetcom-encryption-network/`, 2017.

[91] D Wayne Osgood, Janet K Wilson, Patrick M O'Malley, Jerald G Bachman, and Lloyd D Johnston. Routine activities and individual deviant behavior. *American Sociological Review*, pages 635–655, 1996.

[92] Masarah Paquet-Clouston, Bernhard Haslhofer, and Benoit Dupont. Ransomware payments in the Bitcoin ecosystem. *Journal of Cybersecurity*, 5(1), 2019.

[93] Sergio Pastrana, Alice Hutchings, Daniel Thomas, and Juan Tapiador. Measuring EWhoring. In *Proceedings of the Internet Measurement Conference*, pages 463–477, 2019.

[94] Smith R and Christou ML. Extracting value from their environment: some observations on pimping and prostitution as entrepreneurship. *Journal of Small Business & Entrepreneurship*, 22(1):69–84, 2009.

[95] William Ralston. A dying man, a therapist and the ransom raid that shook the world. `https://www.wired.co.uk/article/finland-mental-health-data-breach-vastaamo`, 2020.

[96] Lee Rondganger. Internet fraudsters swipe thousands from NGO account. `https://www.security.co.za/news/5618`, 2007.

[97] 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.

[98] Molly Sauter. *The coming swarm: DDOS actions, hacktivism, and civil disobedience on the Internet*. Bloomsbury Publishing USA, 2014.

[99] Thomas C Schelling. What is the business of organized crime? *Journal of Public Law*, 20(1):71–84, 1971.

[100] Carl Shapiro and Hal R Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard, 1998.

[101] Susan Leigh Star. The ethnography of infrastructure. *American Behavioral Scientist*, 43(3):377–391, 1999.

[102] Susan Leigh Star and Geoffrey C Bowker. How to infrastructure. *Handbook of New Media: Social Shaping and Social Consequences of ICTs*, pages 230–245, 2006.

[103] Laube Stefan and Rainer Böhme. Strategic aspects of cyber risk information sharing. *ACM Computing Surveys*, 50(5), 2017.

[104] Nicolas P Suzor, Sarah Myers West, Tarleton Gillespie, and Jillian York. Guiding principles for the future of content moderation: Four scholars and advocates in conversation [breakout session output]. *All Things in Moderation*, 2017.

[105] Kan-ichiro Suzuki, Sang-Hoon Kim, and Zong-Tae Bae. Entrepreneurship in Japan and Silicon Valley: a comparative study. *Technovation*, 22(10):595–606, 2002.

[106] Gresham M Sykes and David Matza. Techniques of neutralization: A theory of delinquency. *American Sociological Review*, 22(6):664–670, 1957.

[107] Samaneh Tajalizadehkhoob, Hadi Asghari, Carlos Gañán, and Michel van Eeten. Why them? Extracting intelligence about target selection from Zeus financial malware. In *Workshop on the Economics of Information Security*, 2014.

[108] Jamie Tarabay. How a dated cyber-attack brought a stock exchange to its knees. `https://www.bloomberg.com/news/articles/2021-02-04/how-a-dated-cyber-attack-brought-a-stock-exchange-to-its-knees`, 2021.

[109] Aleksi Teivainen. Hacked Finnish psychotherapy service provider declared bankrupt. `https://www.helsinkitimes.fi/finland/finland-news/domestic/18704-hacked-finnish-psychotherapy-service-provider-declared-bankrupt.html`, 2021.

[110] Daniel R Thomas, Richard Clayton, and Alastair R Beresford. 1000 days of UDP amplification DDoS attacks. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*, pages 79–84. IEEE, 2017.

[111] Anna Tims. 'Sim swap' gives fraudsters access-all-areas via your mobile phone. `https://www.theguardian.com/money/2015/sep/26/sim-swap-fraud-mobile-phone-vodafone-customer`, 2015.

[112] Dan Turkel. Hackers are now offering 'customer support' to the victims they extort money from. `https://www.businessinsider.com/ransomware-writers-offer-customer-support-to-victims-2016-1`, 2016.

[113] UK Finance. Fraud – the facts 2020. `https://www.ukfinance.org.uk/system/files/Fraud-The-Facts-2020-FINAL-ONLINE-11-June.pdf`, 2020.

[114] U.S. Attorney's Office Southern District of New York. Charging Documents: U.S. V. Nikita Kuzmin, U.S. V. Mihai Ionut Paunescu, And U.S. V. Deniss Calovskis. `https://www.justice.gov/usao-sdny/pr/charging-documents-us-v-nikita-kuzmin-us-v-mihai-ionut-paunescu-and-us-v-deniss`, 2013.

[115] U.S. Attorney's Office Southern District of New York. Indictment & Supporting Documents: U.S. V. Liberty Reserve, Et Al. `https://www.justice.gov/usao-sdny/pr/indictment-supporting-documents-us-v-liberty-reserve-et-al`, 2013.

[116] US District Court for the Southern District of California. Case 3.21-mj-01948-MSB. `https://www.documentcloud.org/documents/20799201-operation-trojan-shield-court-record`, 2021.

[117] Lisa Vaas. Two men busted for hijacking victims' phones and email accounts. `https://nakedsecurity.sophos.com/2019/11/18/two-men-busted-for-hijacking-victims-phones-and-email-accounts/`, 2019.

[118] Sophie van der Zee, Richard Clayton, and Ross Anderson. The gift of the gab: Are rental scammers skilled at the art of persuasion? *arXiv:1911.08253*, 2019.

[119] Kim A Wagner. *Thuggee: Banditry and the British in Early Nineteenth-Century India*. Cambridge Imperial & Post Colonial Studies, 2007.

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

[121] David Wall. Dis-organised crime: Towards a distributed model of the organization of cybercrime. *European Review of Organised Crime*, 2(2):71–90, 2015.

[122] David S Wall and Matthew L Williams. Policing cybercrime: networked and social media technologies and the challenges for policing, 2013.

[123] Imogen West-Knights. The joy of 'lawyer cat' is that it teaches us nothing – it's just very funny. https://www.theguardian.com/commentisfree/2021/feb/11/lawyer-cat-funny-texas-rod-ponton-judge, 2021.

[124] Monica T Whitty. Anatomy of the online dating romance scam. *Security Journal*, 28:443–455, 2015.

[125] Jeff Yan. From Sicilian mafia to Chinese "scam villages". *arXiv:1905.03108*, 2019.