SIZING CYBERCRIME: INCIDENTS AND ACCIDENTS, HINTS AND ALLEGATIONS

Stephen Cobb
ESET, USA

Email stephen.cobb@eset.com

ABSTRACT

Cybercrime certainly feels like a major threat to the security of networked systems upon which so much of daily life depends in the world today. Criminals routinely use digital networks to steal data, defraud companies and consumers, and disrupt normal business operation in both public and private sectors. But just how big a threat is cybercrime? For a problem long characterized as both huge and existential by politicians and industry pundits, cybercrime has largely gone unmeasured, if 'measure' is taken to mean 'ascertain the size of the problem using sound scientific methodology'. In this paper we review the cybercrime literature, both commercial and academic, for answers as to why there is a lack of reliable, consistent, longitudinal data on the size and scope of the cybercrime problem. The need for such data is discussed and suggestions for future research are advanced, as well as warnings about excessive reliance on surveys of dubious methodology.

INTRODUCTION

('Hello, Is there anybody in there?' – Pink Floyd)

In 2013, an article appeared in *Communications of the ACM* entitled 'Cybercrime: it's serious, but exactly how serious?' The short answer, and the short version of this paper, is this: we still don't know, but we're working on it. In that article, which highlights the incisive analysis of 'cost of cybercrime' surveys by Florêncio and Herley [1], there is a challenging quote from Ross Anderson [2]: 'Stop wasting money on measuring cybercrime… spend it on the police instead.' Many information security professionals applaud increased funding for efforts to bring cybercriminals to justice – after all, there are other threats against which information systems must be defended besides criminals – but does it make sense to rid ourselves of our current dependency on cybercrime surveys? Hopefully, some historical context and a review of the issues will help us answer this question.

Protecting citizens and their property against harms caused by criminal activity is a basic function of modern society, as is the counting of crimes that occur despite society's efforts to prevent them. Criminology is the science that studies these aspects of human behaviour, and criminologists have a lot to say about measuring crime. The origins of both sociology and criminology are closely related to early efforts at tabulating and analysing 'moral statistics' such as births, deaths, education levels and criminal acts, notably Guerry's research in nineteenth century France [3, 4]. Guerry was the first to map criminal activity and demonstrate the potential to increase our understanding of crime through the analysis of multiple variables, like location and level of education, relative to types of crime. In their 1829 *Statistique Comparée de l'État de l'Instruction et du Nombre des Crimes*, Guerry and Balbi showed that the north of France had the highest levels of both education and property crime [3].

Over the next 100 years, the governments of many countries went through the process of formalizing the collection of social statistics, including recording crimes against persons and property, driven in part by a desire to assess society's progress in efforts to discourage crime. This process involved much debate over what should be recorded – for example, should it be crimes reported or criminals indicted? crimes prosecuted or convictions obtained? By 1930, both the United States and the United Kingdom had settled on the use of 'uniform crime reports' that track specific crimes reported to law enforcement [5]. In the US, eight predatory, common-law classifications were selected: the four 'violent crimes' of homicide, forcible rape, robbery, and aggravated assault; plus the four 'property crimes' of arson, burglary, larceny ($50 and over), and motor vehicle theft [6].

Setting aside the extensive debates over the accuracy and validity of these official records, there is strong consensus that they conceal the 'dark figure' of crime, defined as: 'occurrences that by some criteria are called crime yet that are not registered in the statistics of whatever agency was the source of the data used' [7]. Some crimes that are reported are not recorded, and some crimes are just not reported, for a variety of reasons, starting with ignorance of the crime. As Kabay put it, 'a certain unknown number of crimes of all kinds are undetected' [8]. Even when known, crimes may not be reported for reasons such as: fear of retribution or humiliation; issues of complicity; or a perception that reporting is pointless. In the case of commercial victims, non-reporting may be due to fear of brand damage, exposure to liability, and possible loss of business. (As for reporting computer crimes, the *CSI*/FBI Computer Crime and Security Survey tracked four frequently cited reasons for not reporting to law enforcement: 'negative publicity, competitors would use to advantage, unaware that could report, and civil remedy seemed best' [9].)

One approach to measuring crimes not reported to the authorities is the victim survey. By 1982, both the US and the UK had instituted regular surveys that ask households about their experience of victimization, known as the National Crime Victimization Survey (NCVS) and the British Crime Survey (BCS) respectively. Historically, criminologists were slow to study businesses as crime victims [10], and when computer-related crime began to emerge as a serious issue for businesses, governments were slow to measure it. This led to the emergence of cybercrime surveys conducted by commercial and non-governmental entities, some of whom could be said to have their own agendas [11].

WHY MEASURE CRIME?

('But there's gonna be a meter on your bed' – Leonard Cohen)

The reasons for measuring crime are long standing, numerous and compelling, as are the functions within society that such measurements may assist. One can posit at least nine different 'markets' that seek crime metrics for a variety of reasons, as suggested in Figure 1 (while 'Researchers' could be considered a tenth 'market', we can also assume they serve the other nine).

The motives for measuring crime are also numerous. The following list is adapted from a report from the 2010 Oxford Internet Institute forum: Mapping and Measuring Cybercrime [12]:

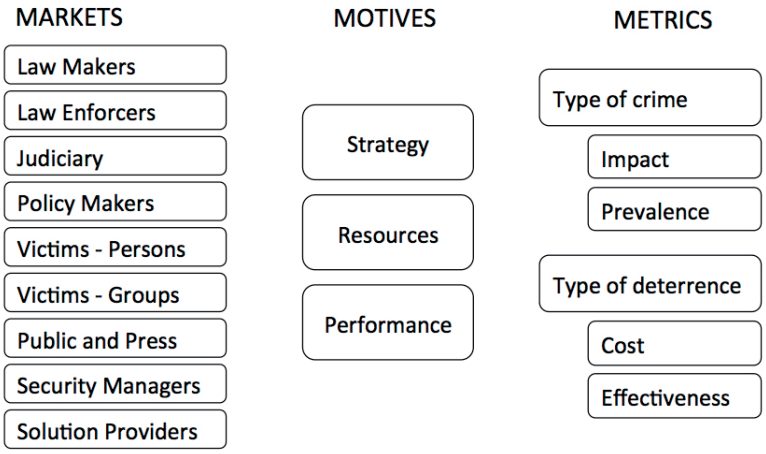

Figure 1: The markets and motives for crime metrics.

- Informing crime reduction initiatives
- Enhancing local and national responses to crime
- Identifying gaps in response to crime
- Providing intelligence and risk assessment
- Identifying preventative measures
- Educating and informing the public
- Identifying areas for further research.

While the above were listed in the context of cybercrime, they apply to all forms of crime, and can be grouped into three categories: strategies, resources and performance. For example, security strategy for banks is informed by knowing that 3,961 bank robberies were reported to the US Federal Bureau of Investigation (FBI) in 2014, and that bank robbery is more likely to occur on Fridays, in commercial districts, and can be fatal (as it was for 10 perpetrators that year) [13]. Such statistics inform all of the functions listed above, particularly when they are compiled every year to create longitudinal data from which additional conclusions can be drawn, and about which useful questions can be asked, such as: why is robbing banks a less popular crime in America today than it was 10 years ago, when there were 7,720 robberies? (In the 1990s, the FBI logged, on average, more than 8,000 bank robberies per year.)

When it comes to cybercrime, there is no official longitudinal collection of statistics, so we are left to make what sense we can of isolated data points. For example, it has been *estimated* that the criminals who stole payment card data from *Target* 'earned' $53.7 million in the first two months of selling it on the black market [14]; but we *know* that the total amount stolen in bank robberies in 2011 was just over $38 million [15]. You don't need to be a criminologist to see that risking death for a crime that yields, on average, less than eight grand, is far less appealing if you can steal and fence data at scale from the relative safety of a comfy chair.

So, while there appears to be a growing consensus that the misuse and abuse of computers and their network connections with criminal intent has become an increasingly serious problem for many businesses, consumers and governments [16, 17], there is still no official barometer to tell us if cybercrime is really on the increase, and if so, at what rate.

Only once has the US government attempted a comprehensive assessment of the impact of cybercrime on business [18]; that was in 2006 and the National Criminal Justice Reference Service (NCJR) recently stated: 'At this time, we do not have any information about any additional reports on this topic becoming available in the future' [19]. This situation is frustrating to numerous constituencies, from the judiciary – 'We need cybercrime metrics to track the damage it inflicts and to develop strategies for dealing with it' [20] – to security managers: 'Don't we owe it to ourselves and to the world to figure out how well we're really doing, instead of leaving it to gut feelings and anecdotal evidence?' [21].

DEFINITIONS

('I can barely define the shape of this moment in time' – Pink Floyd)

So far this paper has glossed over one obvious problem with measuring cybercrime: how to define it. For the purposes of this paper, cybercrime is taken to mean: 'crimes in which computer networks are the target or a substantial tool' [22]. This definition preserves the spirit of one of the earliest definitions of computer crime: a crime 'in which a computer was directly and significantly instrumental' [23]. General objections to 'cybercrime' are noted; however, as Wall observes, '"cyberspace crime" would have been a more accurate descriptor [but] the term "cybercrime" prevails as the accepted term' [24].

The weakness of cybercrime as a categorization, ably explicated by Brenner [20], is also acknowledged, as is the observation made by Anderson *et al.* that: 'the boundary between traditional crime and cybercrime remains fluid' [25]. The terminology of crime also needs a word or phrase to denote crime that can be perpetrated without any digital elements, such as 'physical crime' or the term used in this paper, 'traditional crime'. That seems to work, although we are grateful to Montoya *et al.* for pointing out that it probably has a limited shelf life [26]. We generally eschew 'meatspace crime', although it can be a dramatic counterpart to 'cyberspace crime'[1].

[1] Meatspace appears to originate in Gibson's Neuromancer (1984), entering the OED in 2001.

STUDIES AND SURVEYS: A CRITICAL HISTORY

('Watching the world wake up to history' – Jesus Jones)

The first book on computer crime appears to be that by McKnight in 1973[2]. Parker's *Crime by Computer* appeared in 1976, based on the collection of case studies begun by the Stanford Research Institute (SRI). The first serious review of computer crime studies seems to be 'A Survey of Computer Crime Studies', published in the *Computer/Law Journal* in 1980, by Taber [23].

A shaky start

Taber dismissed McKnight's book as 'too uncritically accepting of media computer "horror stories" to be of any value' and put his finger on a phenomenon now familiar in our industry: 'once in print a computer crime story tends to live forever'. Taber reviewed four source documents, of which a GAO study earned the highest marks. Parker's work at SRI came in for the most criticism, in part because Taber was not happy with the conflation of computer crime and 'computer abuse', the latter concept apparently originating with SRI, probably because they began tracking numerous computer-involved acts of malfeasance before they were declared criminal. Computer abuse was defined as: 'an intentional act in which one or more victims suffered, or could have suffered, a loss and one or more perpetrators made, or could have made, a gain. The incident must be associated with computer technology or its use' [27].

This definition is too far distant from 'crime' for Taber and his critique of the SRI study crystallized tensions that persist to this day: do we count only those adverse events involving computers that constitute criminal acts in which there were quantified and verified losses, or do we track all incidents that detract from the smooth operation of information systems?

As suggested earlier in Figure 1, there are multiple constituencies for data pertaining to computer abuse, from security professionals charged with defending systems, to law enforcement officials responsible for prosecuting lawbreakers, and lawmakers attempting to maintain a state of law and order in which the benefits of digital technology can be enjoyed. Other interested parties include: vendors of security solutions seeking to educate the market for their products; society at large, which is the presumed beneficiary of digital technology's capabilities; and society's self-appointed watchdogs, the media.

Taber was prescient in pointing out that none of these 'markets' are well served by poorly defined data, with the possible exception of less scrupulous vendors and the media, the latter apparently thriving on confusion when it comes to reporting anything computer abuse-related. Taber argued that none of the studies he examined substantiated the authors' claims that computer crime was a problem, declaring '"Computer crime" is a media creature, largely fed by computer security industry press releases' and decrying the drafting of computer crime laws 'justified by such inadequate research' [23].

[2] Ironically, Computer Crime by G. McKnight is currently offered as a free download by a shady streaming content site.

The not so golden age of computer crime surveys

Regardless of what Taber thought, Parker had no illusions about computer crime studies and would later join those calling for caution in the use of computer crime statistics, which proliferated towards the end of the last century. In 1996, a company called *Computer Security Institute* (*CSI*) began conducting a 'Computer Crime and Security Survey' with help from the FBI [328]. *CSI* described itself as a membership organization for information security practitioners and about 4,800 members were contacted by mail for the first survey; 428 responses were received, from people with job titles ranging from corporate information security manager to data security officer and senior systems analyst. The organizations surveyed included corporations, financial institutions, government agencies and universities.

Despite some flags raised by the research methodology, like the low response rate and the fact that responses were anonymous, this first *CSI*/FBI survey was big news. It was cited at length in testimony to the US Senate by its editor, Richard Power, who presented several observations and recommendations that most information security professionals in 1996 heartily endorsed: the problem of computer crime is 'only getting worse'; there is an 'insufficient level of commitment to information security'; there is a need for 'in-depth, periodic risk analysis' and developing 'strong, enforceable policies on a broad range of information security issues'; there is a need for 'security awareness for users'; and 'a great need for an emphasis on information security in computer science curriculum and on computer ethics as a critical aspect of good citizenship' [28].

For validation of the survey's results, Powers looked beyond its methodological shortcomings and cited other studies that appear to show similar results, together with reports of incidents that exemplify the survey's findings. *CSI* reports were produced annually until 2010 when the project went dark, but the template that Powers created endures to this day: conduct computer crime survey, validate by cross-reference, make recommendations based on results, garner attention, and boost security awareness. Unfortunately, the template has serious flaws. Consider comments from Schneier in 2001:

'The results are not statistically meaningful by any stretch of the imagination – they're based on about 500 survey responses each year… This data is not statistically rigorous, and should be viewed as suspect for several reasons. First, it's based on the database of information security professionals that the *CSI* has (3,900 people), self-selected by the 14% who bothered to respond. (The people responding are probably more knowledgeable than the average sysadmin, and the companies they work for more aware of the threats…) Second, the data is not necessarily accurate, but is only the best recollections of the respondents. And third, most hacks still go unnoticed; the data only represents what the respondents actually noticed.' [29]

Despite this, Schneier also says, 'but it is the most interesting data on real-world computer and network security that we have', thereby crystallizing the computer crime survey dilemma: we're so short of metrics, we'll take what we can get, even when the purveyor of those metrics warns us of their limitations. For example, Power had few illusions about the *CSI* data and took pains to include caveats in every edition of

the survey, including warnings from Parker in 1999 and those of Schneier, cited above, in 2002.

New century, new fears

Clearly, by the end of the last century, information security professionals were on notice that computer crime surveys might not accurately reflect reality. A white paper by Kabay in 1998 had spelled out the statistical realities, including the ascertainment problem; this formed the basis of a chapter on the topic in *Wiley*'s *Computer Security Handbook* [8]. Yet, the surveys kept coming. By 2003, Ryan and Jefferson felt compelled to publish 'The Use, Misuse, and Abuse of Statistics in Information Security Research', in which they analysed 14 different surveys [30]. Twelve were found to have compounded erroneous extrapolations of data by failing to limit responses to one per company. The resulting reports were sometimes presented as being representative of company experience, even though a statement that 'two thirds of companies experienced cybercrime committed by an employee' may be entirely inaccurate if more than one response per company is allowed.

By 2005, the range of companies reporting computer crime metrics included security vendors like *Symantec* and the big consulting firms that became *EY*, *PwC*, *Deloitte* and *KPMG*. These private sector efforts were joined by some government initiatives, such as the annual reports produced by the FBI's Internet Crime Complaint Center (IC3). In 2000, IC3 started reporting on many different Internet-related complaints, mainly from US citizens [31]. In 2014, the organization received 269,422 complaints from the public and estimated the associated dollar losses at over $800 million. While IC3 performs a very valuable role triaging and referring for investigation all manner of Internet-related rip-offs, as well as alerting consumer organizations to the latest scams, the data it reports has limited value as a precise measure of cybercrime. As Viega pointed out, reporters are self-selecting and validation of whether or not a crime was actually committed is weak [21]. Furthermore, the criteria by which reports from victims in other countries are included are unclear. As for the calculation of losses incurred, these depend too heavily on unchecked victim estimates. As Ryan and Jefferson noted, relying on such estimates introduces considerable potential for error [30], a phenomenon that Florêncio and Herley would later take great pains to explain [1].

Another series of surveys debuted in 2004, associated with the Carnegie Mellon University Software Engineering Institute's CERT Coordination Center (CUC). The E-Crime Watch Survey was conducted online by *CSO Magazine* with help from both CUC and the US Secret Service [32]. This survey was run again in 2005 and subsequently, with *Microsoft* involvement, in 2006 and 2007[3]. It apparently disappeared in 2008 but re-emerged in 2010/11/12 as the CyberSecurity Watch Survey (with *Deloitte*). This morphed into the US State of Cybercrime Survey in 2013 and appeared under the same name in 2014 but with *Pricewaterhouse Coopers* (*PwC*) instead of *Deloitte* [34]. The number of respondents to these surveys ranged between 500 and 1,000.

While the original 'Watch surveys' were very upfront about methodology and limitations, the 2014 'State of Cybercrime' report exemplifies both the selective presentation of results and inconsistency of terminology that has plagued commercially sponsored cybercrime surveys, for example mixing and matching terms like: crime, attack, breach and incident. While billed as a survey of cybercrime in the United States, the *PwC* report offers no definition of what constitutes a cybercrime and uses the terms 'incident' and 'threat' more than 'crime' or 'cybercrime'. For example, the report states that 72% of respondents think outsiders like hackers are the source of 'cybersecurity incidents', and contrasts this with 'insider events' reported by only 28% of respondents. The document goes on to note that one third of respondents think 'insider crimes are more costly or damaging than incidents perpetrated by outsiders'. This same statistic appears in a *PwC* report on the global state of information security [35]. Unlike the US report, the global report declares insiders to be 'the most-cited culprits of cybercrime', but the data appears to come from a table titled 'Sources of security incidents'. While mixing crimes and incidents undermines efforts to use the data for security decision-making, a deeper issue is the lack of access to the actual survey data, or even a full set of results. *PwC* only releases 'Key Findings' and asserts that 'the full report is proprietary and not available to the public' [36].

The government steps in

In 2006, the US federal government's Bureau of Justice Statistics (BJS) took a stab at measuring the cost of cybercrime to businesses and the results were relatively impressive. More than 36,000 businesses were contacted for the National Computer Security Survey, a stratified, random sample designed to produce national and industry-level estimates [37]. Over 8,000 businesses responded, an overall response rate of 23%. By comparison, the *CSI*/FBI survey averaged 500 responses and the response rate never exceeded 15%. Nevertheless, the BJS statistician made it clear that even this level of response was 'not sufficient to support national or industry-level estimates'.

While the report turned up 22 million incidents of cybercrime in 2005, the vast majority – 20 million – were defined as *other computer security incidents*, primarily spyware, adware, phishing and spoofing. As to the cost of this activity, 91% of the responding businesses sustained monetary loss and/or system downtime, with the monetary losses pegged at $867 million. Sadly, the US government never repeated this study, which would have at least given us multiple snapshots in time to compare (requests to the BJS for more recent data on cybercrime's impact on businesses are currently referred to the 'Key Findings' report by *PwC* [19]). Not that all is well with other government crime reporting. For example, units of measure are apt to change, and there are gaps: for example, the BJS figures for identity theft for 200–2010 are by households affected, but 2012 data are by persons; and the FBI lost some of the bank robbery data pertaining to 2012 to 2014 [15].

Practical impossibilities?

One possible explanation for the DoJ's decision not to repeat the 2005 study is that, as Florêncio and Herley explain, large stratified random samples are necessary if you want to generalize results [1]. As with personal wealth, the concentration of cybercrime losses tends to be unevenly distributed across the population and so representative

[3] Missing reports are referenced by media, e.g. 2007 in Information Week [33].

sampling gives an unrepresentative estimate. When losses are limited to a small group of companies within survey samples, this tends to amplify the problems caused by small sample sizes and low response rates with respect to reported losses. Consequently, 'outliers can cause catastrophic errors'. Florêncio and Herley helpfully catalogue the ways in which erroneous outliers can occur, notably through a lack of input validation, an ironic weakness in the case of computer crime surveys (classically demonstrated by the gap between the number of sex partners reported by men and women [1]).

Another practical impossibility would appear to be educating the various 'markets' for cybercrime statistics on their appropriate use. While many of us security professionals should know better, too often we find it hard to avoid citing statistics to make a point, but all too easy to avoid making the appropriate disclaimers; at times it is as though we are co-dependent with the entities that publish the numbers.

Consider the annual *Verizon* Data Breach Investigation Report (DBIR). The first DBIR appeared in 2008, and DBIRs have been widely quoted every year since. From the very first edition, this report included a disclaimer, while making a habit of announcing numbers in a manner that conveys a confidence not necessarily supported by the data. Take two examples from the 2015 report: mobile malware and the cost of breaches [38]. Under the catchy headline 'I Got 99 Problems and Mobile Malware Isn't Even 1% of Them', the report gives the impression that mobile devices are not a serious breach vector. Yet an infection rate of 0.03% is cited for 'higher-grade' malicious code on smartphones, which translates to over 33,000 seriously compromised devices in the workplace, based on 140 million Americans at work, and 80% of them using their personal devices for work [39]. One is reminded that the *Verizon* data set is limited to a subset of known breaches. As for the cost of breaches, the 2015 DBIR declares that this is $0.58 per record, a number strongly disputed by the *Ponemon Institute* [40], which puts the figure closer to $200, in a classic example of the 'orders of magnitude' problem cited by Hyman [2].

Damage assessment as a crime metric

Regardless of survey results, it is easy to feel that the prevalence and impact of cybercrime are currently higher than ever, yet the Internet appears to keep growing, in terms of users, nodes, traffic, applications, and so on. But is it growing as fast as it would if there were less cybercrime? One useful possible measure of cybercrime might be the extent to which it discourages adoption of Internet technology; and there is strong evidence that 'a substantial fraction of cybercrime's overall costs to society can be traced to indirect opportunity costs, resulting from unused online services' [17, 41]. In a recent study, Riek *et al.* used European survey data to confirm the 'negative impact of perceived risk of cybercrime on the use of all three online service categories' (these were banking, shopping, and social networking) [17]. The data came from the European Commission's annual Eurobarometer Cyber Security Survey, the third edition of which was published this year, compiling responses from more than 27,000 respondents in 28 countries.

While the EU publishes the top line findings of the Cyber Security Survey, including high levels of concern about online security, the raw data is also made available, enabling secondary analysis of the type performed by Riek *et al*. They analysed the 2013 data, which clearly showed that perceived risk of cybercrime deters use of online services, increasing 'avoidance intention'. The ability to repeat this and other studies from a full and solid set of data offers considerable potential to answer questions such as: when does cybercrime reach levels that seriously impact the digital economy?

DISCUSSION

('Is it a crime?' – Sade)

Clearly, accurate measurement of the prevalence and cost of any criminal activity is extremely difficult, and subject to far greater margins of error than are conveyed by the most widely read and quoted cybercrime surveys. The most comprehensive and potentially reliable survey conducted in the United States required extensive resources and was arguably a project that only the federal government could undertake, but the government has no plans to repeat the study. What is worse, the government now points people to exactly the kind of survey that numerous experts have convincingly invalidated. Despite this, concern about cybercrime is clearly valid; just ask any citizen who has suffered identity theft as a result of a security breach, or any security manager who is working to prevent such breaches. In light of this, what useful observations can be made based on our review of the cybercrime measurement?

The right answer might be: we don't know

Good security researchers know the limits of their knowledge and confidently assert that limit when asked questions for which there is no good answer. At least that is the way it should be. Unfortunately, within the security industry, this aspiration is not always appreciated by the folks in sales and marketing and PR, not to mention the media and the general public. There is a big demand, and widespread exposure awaits, for numbers that give dimension to the great unknowns, such as: is cybercrime increasing? What is the most common form of attack? Who hacked *Sony*?

That last question is a tricky one, given that there have been numerous hackings of cyber assets that have *Sony* in the name, but it may point the way, not to better metrics, but to better use of the metrics we do have. Many security researchers, notably those involved with malware investigations, appear to be comfortable stating that attribution is very difficult, if not impossible. That doesn't downplay the threat or make it less real, nor does it change the security strategy required to defend against it, so perhaps we should more frequently assert that measuring cybercrime is very difficult, if not impossible.

X amount of crime = cost of doing business?

Here's another telling question: how much did the *Target* breach cost the company? Answer: not as much as light-fingered employees. Retail establishments are routinely victimized by criminals. As one of the world's oldest trades, retailing has had centuries in which to come to terms with the reality of things going missing, falling off the back of a cart, being billed but never delivered, lifted by customers and employees, and so on. There is even a name for this: shrinkage, which encompasses shoplifting, employee theft, vendor fraud, accounting errors and process failures [42].

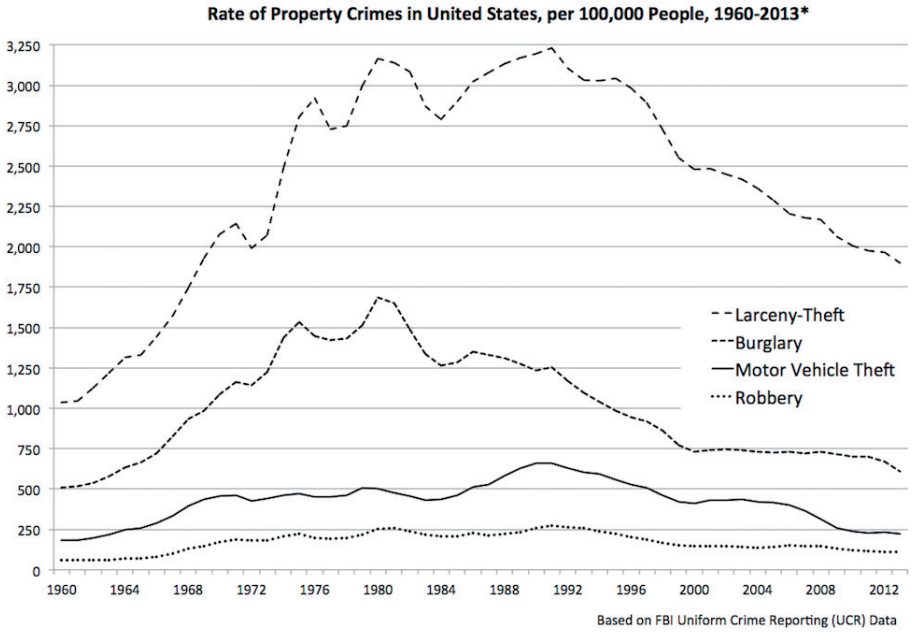

Figure 2: Property crime rates in the US.

Shrinkage is often expressed as a percentage of retail revenue, for example 1.48% in the US, somewhat less in the UK. In the UK, employees are thought to account for 15.9% of shrinkage. In the US that figure is 42.9%. In monetary terms, US retailers lost $43 billion to employees in 2013. Even a conservative calculation of employee theft at *Target*, which apparently has a good reputation for shrinkage control [43], puts the annual loss at over $400 million in 2013, roughly twice the $200 million figure cited for the cost of the 2013 data breach.

For cybersecurity folks there is plenty of schadenfreude to go around here. For a start, retail security experts are by no means happy with the way shrinkage is calculated, arguing about definitions, methodology, and so on [44]. More pertinently, retail provides an example of how companies can function successfully while enduring X amount of criminal activity. Security researchers should at least consider whether this could be a model for coping with some forms of cybercrime.

CONCLUSIONS

The simplest explanation for why we lack reliable, consistent, longitudinal data on the size and scope of the cybercrime problem is that such data is hard to get. The temptation to cut corners is great, as is the temptation to oversell the reliability of the results of surveys that are methodologically weak. For all the talk about a national emergency due to 'the increasing prevalence and severity of malicious cyber-enabled activities' [45], the US government appears to have given up on measuring cybercrime and is even making a mess of historically helpful data [15]. This is a shame, because measuring crime can produce valuable knowledge. For example, when cybersecurity professionals take a moment to look beyond botnets and data breaches to the world of traditional crime there appear to be some trends worth smiling about. The crime wave that rose up in the 1960s and persisted through the 1980s actually crested in the early 1990s and has been receding ever since (see Figure 2).

At first glance, this pattern, prosaically dubbed 'the crime drop' by criminologists, would seem to offer hope to beleaguered consumers and CISOs alike. Although the causes of this phenomenon, also seen in the UK and other countries, are keenly debated, it could mean that crime rates *can* go down, which means we may not be fighting cybercrime forever. Sadly, we do not have the data to say this with any certainty. Criminologists are slowly beginning to realize that the crime drop shown in Figure 2 might have something to do with cybercrime[4]. Indeed, the crime drop may be a form of crime displacement, a phenomenon well documented in the literature of traditional crime but not applied to cybercrime until relatively recently [49]. There are five types of crime displacement: changing location, changing timing, choosing alternative targets, applying different tactics, and trying a different type of crime [50]. Could criminals be trying cybercrime instead of 'traditional crime'? Exploring this possibility could be a productive focus of future cybercrime research. Another promising direction is indicated by Riek *et al.* who note 'research on online service avoidance as a response to perceived risk of cybercrime is rare and isolated' [17]. If such research revealed that fear of cybercrime was hurting corporate profits then we might see greater attention paid to combating cybercrime.

So, was Anderson correct, should we shift the cybercrime survey budget to law enforcement? In the US, where the US government appears to be spending very little money on gathering reliable cybercrime metrics, the question is somewhat academic, and a partnership between private industry and academia might be the way forward if we still have the heart to try and measure this problem, and the willpower to hold our biases in check (while concurrently lobbying our elected representatives to increase funding for sensible cybercrime deterrence measures).

[4] Some clue here: [46]; no clue here: [47]; lots of clues here: [48].

REFERENCES

[1] Florêncio, D.; Herley, C. (2013). Sex, lies and cyber-crime surveys. Economics of information security and privacy III, New York: Springer 35–53. http://research.microsoft.com/pubs/149886/SexLiesandCybercrimeSurveys.pdf.

[2] Quoted in Hyman, P. (2013). Cybercrime: it's serious, but exactly how serious? Communications of the ACM, 56(3), 18–20.

[3] Friendly, M. (2007; 2008). A.-M. Guerry's "Moral Statistics of France": Challenges for Multivariable Spatial Analysis. Statistical Science, 22 (3): 368–399.

[4] Guerry, A.; Whitt, H.; Reinking, V. (2002). A Translation of Andre-Michel Guerry's Essay on the Moral Statistics of France (1883): a sociological report to the French Academy of Science (Vol. 26). Edwin Mellen Press.

[5] Maguire, M. (2007) Crime data and statistics. in Maquire, M.; Morgan, R.; Reiner, R. (Eds.) The Oxford handbook of criminology (4th Edition), Oxford: Oxford University Press, 241–301.

[6] FBI Uniform Crime Reporting (UCR) Program. https://www.fbi.gov/about-us/cjis/ucr/ucr.

[7] Biderman, A.; Reiss, A. (1967). On Exploring the "Dark Figure" of Crime. Annals of the American Academy of Political and Social Science, 374 (1): 1–15.

[8] Kabay, M. (2014). Understanding studies and surveys of computer crime. In Bosworth, S.; Kabay, M.; Whyne, E. (Eds.) Computer Security Handbook. Hoboken, New Jersey: John Wiley & Sons, Inc. 10.1–10.12.

[9] Power, R. (2000), 2000 CSI/FBI computer crime and security survey. COMPUT SECUR J, 16(2), 33–49.

[10] Gill, M. (Ed.). Crime at Work: Studies in Security and Crime Prevention. Leicester: Perpetuity Press.

[11] Anderson, R.; Böhme, R.; Clayton, R.; Moore, T. (2008). Security economics and the internal market. Study commissioned by ENISA.

[12] Fafinski, S.; Dutton, W.; Margetts, H. (2010). Mapping and Measuring Cybercrime. OII Working Paper No. 18, SSRN Working Paper Series.

[13] FBI Bank Crime Reports. https://www.fbi.gov/stats-services/publications.

[14] Krebs, B. (2014). The Target Breach, By the Numbers. Krebs on Security. http://krebsonsecurity.com/2014/05/the-target-breach-by-the-numbers.

[15] Due to 'an upgrade in our case management system' the FBI is not able to provide complete bank robbery data for 2012, 2013, or 2014. Author's correspondence with FBI, July, 2015.

[16] Hernandez-Castro, J.; Boiten, E. (2014). Cybercrime prevalence and impact in the UK. Computer Fraud & Security, 2014 (2), 5–8.

[17] Riek, M.; Böhme, R.; Moore, T. (2015). Measuring the Influence of Perceived Cybercrime Risk on Online Service Avoidance. IEEE Transactions on Dependable and Secure Computing PP, 99, 1–1.

[18] US Department of Justice (2008). Cybercrime against businesses, 2005. http://www.bjs.gov/content/pub/pdf/cb05.pdf.

[19] Author's personal correspondence with National Criminal Justice Reference Service (NCJRS), November, 2014.

[20] Brenner, S. (2004) .Cybercrime Metrics: Old Wine, New Bottles? Virginia Journal of Law & Technology, 9, 13–13.

[21] Viega, J. (2012). Ten years on, how are we doing? (Spoiler alert: We have no clue). IEEE Security & Privacy, (6), 13–16.

[22] Koops, B. (2011). The Internet and its Opportunities for Cybercrime. Tilburg Law School Legal Studies Research Paper Series, No. 9/2011.

[23] Taber, J. (1980). A Survey of Computer Crime Studies 2. Computer Law Journal, 275. http://repository.jmls.edu/jitpl/vol2/iss1/15.

[24] Wall, D. (2008). Cybercrime and the Culture of Fear: Social Science Fiction(s) and the Production of Knowledge about Cybercrime (Revised Feb. 2011). Information, Communication & Society, Vol. 11, No. 6: 861–884.

[25] Anderson, R.; Barton, C.; Böhme, R.; Clayton, R.; van Eeten, M.; Levi, M.; Moore, T.; Savage, S. (2013). Measuring the cost of cybercrime. In Böhme, R. (Ed.) The Economics of Information Security and Privacy. Heidelberg: Springer Berlin, 265–300.

[26] Montoya, L.; Junger, M.; Hartel, P. (2013). How "Digital" is Traditional Crime? IEEE, 31–37.

[27] Belzer, J.; Holzman, A.; Kent, A. (1978). Encyclopedia of Computer Science and Technology: Volume 6 – Computer Selection Criteria to Curriculum Committee on Computer Science (Vol. 6). CRC Press. p.384.

[28] Power, R. (1996). Testimony of Richard G. Power, Editor, Computer Security Institute Before the Permanent Subcommittee on Investigations. US Senate Committee on Governmental Affairs. http://fas.org/irp/congress/1996_hr/s9606051.htm.

[29] Schneier, B. (2001). CSI's Computer Crime and Security Survey. Crypto-Gram. https://www.schneier.com/crypto-gram.

[30] Ryan, J.; Jefferson, T. (2003). The Use, Misuse and Abuse of Statistics in Information Security Research. Proceedings of the 23rd ASEM National Conference. ASEM 15–18 October 2003.

[31] IC3 annual reports from 2001 to 2014. https://www.ic3.gov/media/annualreports.aspx.

[32] Partial list at CERT. http://www.cert.org/insider-threat/research/cybersecurity-watch-survey.cfm?.

[33] http://www.informationweek.com/government/cybersecurity/survey-shows-security-pros-overconfident/d/d-id/1059021?.

[34] PwC (2014). US Cybercrime: Rising Risks, Reduced Readiness – Key Findings from the 2014 US State of Cybercrime Survey. PricewaterhouseCoopers, CERT

Division of the Software Engineering Institute, CSO Magazine, & United States Secret Service.

[35] PwC (2014). Managing cyber risks in an interconnected world – Key findings from the Global State of Information Security Survey 2015, PwC.

[36] Author's personal correspondence with PwC, February 2015.

[37] US Department of Justice. Cybercrime against businesses, 2005. http://www.bjs.gov/content/pub/pdf/cb05.pdf.

[38] Verizon DBIR.

[39] Harris Interactive poll conducted for ESET, 2012. http://www.welivesecurity.com/2012/02/28/sizing-up-the-byod-security-challenge/.

[40] Why Ponemon Institute's Cost of Data Breach Methodology Is Sound and Endures. http://www.ponemon.org/blog/why-ponemon-institute-s-cost-of-data-breach-methodology-is-sound-and-endures.

[41] Riek, M.; Böhme, R.; Moore, T. (2015). Measuring the Influence of Perceived Cybercrime Risk on Online Service Avoidance. IEEE Transactions on Dependable and Secure Computing PP, 99, 1–1.

[42] Checkpoint Systems. The Global Retail Theft Barometer 2013–2014 October 2014. http://www.globalretailtheftbarometer.com/

[43] Beck, A. Effective Retail Loss Prevention: 10 Ways to Keep Shrinkage Low. University of Leicester, 2007. https://www2.le.ac.uk/departments/criminology/people/bna/10WaystoKeepShrinkageLowpdf.

[44] Beck, A.; Peacock, C. (2009). New loss prevention: Redefining shrinkage management. Palgrave Macmillan.

[45] Obama, B. (2015). Blocking the Property of Certain Persons Engaging in Significant Malicious Cyber-Enabled Activities. Executive Order 13694, Federal Register 80 (63): 18077–79. https://www.federalregister.gov/articles/2015/04/02/2015-07788/blocking-the-property-of-certain-persons-engaging-in-significant-malicious-cyber-enabled-activities.

[46] Farrell, G.; Tseloni, A.; Mailley, J.; Tilley, N. (2011). The Crime Drop and the Security Hypothesis. Journal of Research in Crime and Delinquency, 48 (2): 147–175.

[47] Lynch, M. J. (2013). Reexamining Political economy and crime and explaining the crime drop. Journal of Crime and Justice, 36(2), 248–262.

[48] Tcherni, M.; Davies, A.; Lopes, G.; Lizotte, A. (2015). The Dark Figure of Online Property Crime: Is Cyberspace Hiding a Crime Wave? Justice Quarterly http://www.tandfonline.com/doi/abs/10.1080/07418825.2014.994658#.VaXvdJNVhBc.

[49] Cobb, S. (2015). Does crime just hop online. S. Cobb on Security (in press).

[50] Reppetto as cited in Gabor, T. (1990). Crime displacement and situational prevention: toward the development of some principles. Canadian Journal of Criminology, 32 (1): 41.