

A Continuum of Internet-Based Crime: How the Effectiveness of Cybersecurity Policies Varies across Cybercrime Types

Eric Jardine

Research Fellow, CIGI

Chapter to be published in *Research Handbook on Digital Transformations*, edited by

F. Xavier Olleros and Majlinda Zhegu. Edward Elgar, Northampton, MA.

Abstract: Internet-based crime is a looming problem. Some crimes, such as the sale of drugs and guns and the distribution of child abuse imagery, have shifted into the applications, forums and chatrooms of the Internet. Other crimes, such as data breaches and identity theft, ransomware and distributed denial of service attacks, are launched via the infrastructure of the network. Cybersecurity policies designed to counter these activities have predictably different effects across the various types of Internet-based crime. Dark Web indexing and Internet service provider botnet mitigation strategies, for example, affect some forms of cybercrime more than others. In particular, this paper outlines how Internet-based crime varies along a continuum from crime *in* the applications of the Internet to crime *via* the infrastructure of the system. It then shows how cybersecurity strategies have differential effects across the various cybercrime types.

Introduction

The Internet is a wonderful tool. One estimate by the Boston Consulting Group suggests that the Internet could contribute upwards of 4.2 trillion dollars to the world economy in 2016, a figure which is 1 trillion dollars larger than the entire economy of Germany in 2014 (Dean et al., 2012). Another study by the consulting firm McKinsey & Company has found that, across 13 connected nations, the Internet contributes an average of 3.4 percent to national GDP and is often the fastest growing sector of the economy (Pélissié du Rausas et al., 2011). At an individual level, the CIGI/Ipsos Global Survey of Internet Security and Trust (2014) even found that, across over 23,000 Internet users in 24 countries, fully 81 percent of respondents indicate that the Internet is crucial for their economic livelihood.

The Internet is more than just an instrument of economic wealth and prosperity. People are increasingly using it as a vehicle for free expression, social engagement, personal recreation and as a means of accessing scientific knowledge. The CIGI/Ipsos survey is again enlightening. Fully 83 percent of individuals surveyed rank the Internet as being either very or somewhat important to their ability to freely express their political and social ideas. Eighty-five percent of those surveyed see the Internet as core to their social engagement, while 87 percent think it is crucially important for their personal enjoyment. And, as Friesike and Fecher highlight in this volume, 91 percent of respondents maintain that the Internet is central to their ability to access scientific knowledge.

As wonderful a tool as the Internet is, it is also essentially neutral (Jardine, 2015a): what the network of networks does depends on how people choose to use it. Millions of civic and well-meaning people have joined the Internet, but, increasingly, a host of nefarious actors have followed suit. Online crime and criminal behaviour have seen a secular increase over time, as the overall size and level of activity in cyberspace has grown (Jardine, 2015b). Unsurprisingly, and as is the case for most other tools, the Internet can be used in both noble and nefarious ways.

In this chapter, I argue that all forms of Internet-based crime fall along a continuum based upon how criminals actually use the technology of the network. The continuum ranges from crimes that occur on top of the network in various Web-based applications to crimes that use the Internet as a means to transmit criminal misdeeds from one device to another. Many crimes, as we shall see, combine elements of both and fall somewhere in the middle of the range. I argue further that understanding the distinction between the use of the applications of the Internet as a forum for crime and the use of the network as a means of transmitting a crime can have serious implications for policymakers. In short, policy makers are best served by knowing what type of crime they are dealing with when devising and implementing cybersecurity policies.

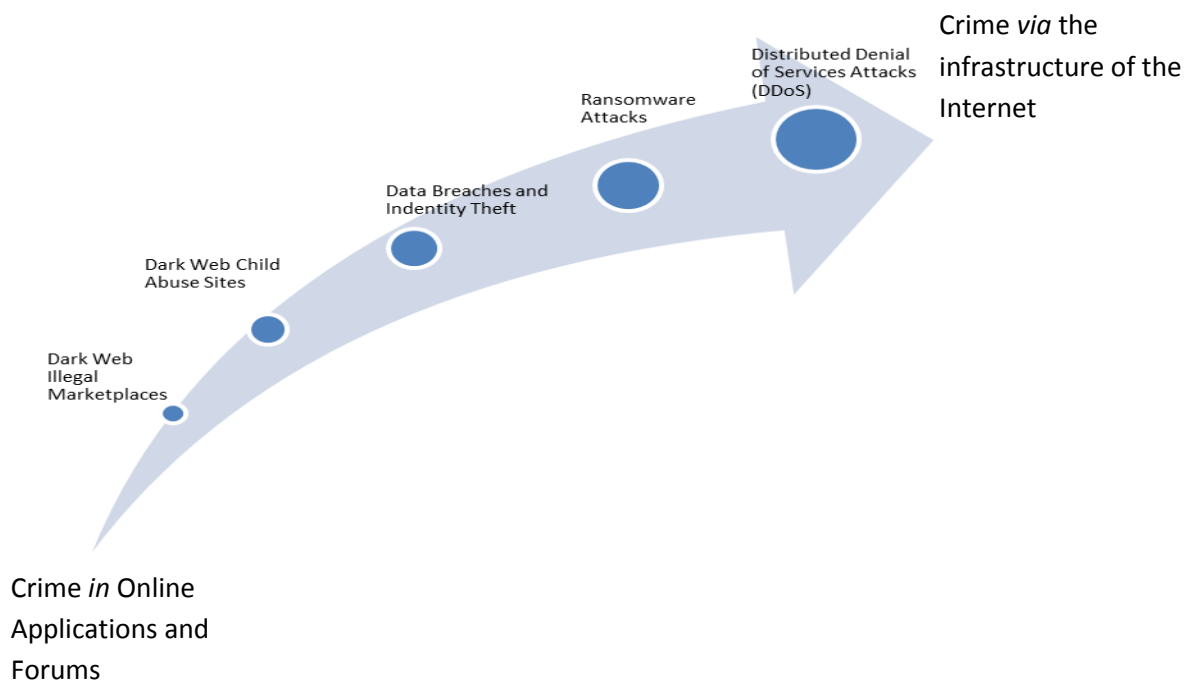
In the next section, I lay out a continuum of Internet-based criminal behaviour that ranges from the occurrence of crime in the applications that ride on top of the Internet to the commission of crimes that occur via the infrastructure of the network of networks. I show further that many Internet-based crimes combine elements of both ends of the continuum and so naturally fall somewhere in the middle. In the

third section, I point to how policies aimed at improving the security of cyberspace have a differential effect on the occurrence of various types of Internet-based crime, depending upon whether the criminal behaviour in question relies mainly on web-based applications or on the Internet as a system via which crime can be committed. In particular, I point to how indexing Dark Web sites can have the largest effect on crime committed in various Web-based applications and a smaller effect on the commission of cybercrime via the infrastructure of the network. I also show how Internet service provider (ISP) botnet mitigation strategies can have a large effect on the commission of crime via the network, but next to no effect on crime that occurs in web-based applications.

Crime and the Internet

To best counter crime in the digital world, policymakers need to understand that not all types of Internet-based crime are the same. Internet-based crime (what we could also call cybercrime, broadly defined) forms a continuum based upon how the criminals makes use of the technology of the system. As shown in Figure 1, one end of the continuum involves crimes that occur in the applications that ride on top of the infrastructure of the Internet. As I discuss in more detail below, illegal online marketplaces fall on this end of the spectrum. On the other end of the continuum are crimes that are committed via the infrastructure of the Internet. As I discuss more below, distributed denial of service (DDoS) attacks that aim to disrupt the operation of a website are probably the best example of a criminal misdeed falling on this end of the continuum. In between the two extremes, there are a number of different crimes, such as child abuse imagery sites, data breaches and identity theft and ransomware attacks, that combine to varying degrees both the use of applications and commission of crime via the Internet's infrastructure.

Figure 1. A Continuum of Internet-Based Crime



The Dark Web

Crime can now easily occur online. Criminals that used to get up to no good offline have widely recognized that the Internet is efficiency enhancing and have shifted the location of their activity. Online criminals now routinely use the applications that ride upon the infrastructure of the Internet as a set of forums for crime and criminal undertakings, essentially replacing the old forums of criminal congregation, such as shady taverns and ill-lit street corners.

Although online crime can occur anywhere on the Internet, a lot of it is undertaken on a part of the network known colloquially as the Dark Web. On the Dark Web, everything is done anonymously. You cannot reach the Dark Web via normal Internet browsers like Google Chrome or Internet Explorer. Instead, you need to use a specially configured browser. The most commonly used system is what is known as The Onion Router, or, as it is commonly referred to, Tor.

Tor allows someone to surf the web anonymously by breaking up the way in which a computer communicates with websites. A typical Internet connection is direct. A person sitting at home or work opens an Internet browser, types in the information that they want to retrieve and then their query is sent via their Internet service provider (ISP) to the website, which returns the desired content. The directness of this approach is why a website's operator knows everything that you view while you are on the site and why the ISP that carries a user's packets of data is able to provide law enforcement with detailed records of all the websites that a person visits, often going back months.

Tor creates anonymity and generates the Dark Web by breaking the signal apart and heavily encrypting a person's web traffic while it is in transit. Rather than communicating directly with a website that a person wants to visit, the Tor browser sends a person's web query to its final destination via a series of three volunteered relay computers. If any individual node was to be co-opted, it could theoretically identify its closest neighbour, but could not link a particular user to the precise content that is being viewed. In order to break the anonymity of the system, law enforcement would need to have control of all three of the computers in a particular relay link to be able to piece together the user and the content that they viewed.

Tor can do more than just allow a user to surf the web anonymously. The Tor network overlay system also includes a series of volunteered servers that can be used to host websites on what are called the Tor Hidden Services. These websites are hosted anonymously, modified anonymously, and, so long as a person knows the website address, accessed anonymously via the Tor browser.

Tor provides anonymity for both users and those wishing to host hidden websites. The anonymity created by Tor provides cover. While that cover can sometimes be good, allowing people living under repressive regimes to exercise their basic political rights to free expression, privacy and access to information (Jardine N.D.), it can also be used to cover up for the commission of crime online (Jardine 2015a). Two examples of crime that used to happen offline and that now occur in Dark Web forums are

illegal marketplaces and child-abuse imagery sites. Exploring each in turn provides a good sense of how crime can occur on top of the network in the applications and forums of the Dark Web.

Illegal Online Marketplaces

In a strict sense, illegal marketplaces are not in any way new. For example, back in the pre-Internet days (and of course to a large extent still), people buy drugs from dealers on the streets. The exchange might happen via cash and be made on a street corner, a public park or a schoolyard. But, at a conceptual level, the physical marketplace, while often transient and informal, is real for a specific moment and time. Other, similar markets did and still do exist for guns, sexual gratification or stolen items and information.

Yet, more and more, these illegal marketplaces are finding a home online. The most famous (or infamous) of these sites was called Silk Road. The website was hosted on the Tor Hidden Services Dark Web. It first emerged in February 2011. It was eventually taken down in 2013, as a result of a Department of Homeland Security (DHS)-led operation called “Marco Polo”. Another iteration of the site was launched only a month later and was, itself, taken down by a law enforcement operation less than a year later (Jardine, 2015a). And again, once Silk Road 2.0 was taken down, another iteration of the site was launched within the hour, although it has failed to reach the same scale as the previous two iterations.

Silk Road 1.0 provides a good example of how a lot of crime has simply migrated online into the applications and forums that ride on top of the Internet’s infrastructure from the offline world. It was a forum for the purchase and sale of all kinds of illegal goods, ranging from guns to drugs. One study found that while there was a fair amount of material for sale on the site, “the four most popular categories are all linked to drugs” (Christin, 2012, 8). Beyond the top four, the study also concluded that 90 percent of the top 10 categories and 80 percent of the top 20 all related to drugs.

Silk Road was very profitable. Ross Ulbricht, who went by the moniker ‘Dread Pirate Roberts’, was behind the first Silk Road. At the Ulbricht trial, the prosecution later alleged that the site had generated upwards of 200 million dollars in illegal drugs sales in just two years (Ax, 2015). As a forum, Silk Road basically put buyers and sellers into contact. It was like an illegal version of Kijiji, Craigslist, or even eBay. The online dealers also pushed a lot of drugs and made a lot of money. One New York seller, Michael Duch, who was charged in the aftermath of the takedown of Silk Road 1.0, provides an interesting snapshot. In a six-month period, Duch sold heroin to some 32,000 individuals to finance his own addiction. Moving huge volumes of highly illegal drugs during these months provided Duch with a huge income, as reportedly he was pulling in between 60,000 to 70,000 dollars per month (Greenberg 2015).

Silk Road changed very little about the occurrence of crime other than its location. Before the Internet, people like Duch would have likely had to turn to selling drugs offline to feed their own addictions (although he claimed that he would not have done so but was lured into dealing by the anonymity of Tor). The Silk Road example clearly shows how crime has migrated online from the physical world into the applications that ride on top of the Internet’s infrastructure.


Child-Abuse Imagery Sites

The production and dissemination of child-abuse imagery obviously long predates the Internet (Bartlett, 2014). Before the Internet, pedophiles would subscribe to child abuse magazines or actively share private photos and literature amongst themselves. The Internet changed the means by which these images were produced and shared, but not the occurrence of the crime itself. Now, pedophiles congregate in Dark Web chat forums and share pictures and videos anonymously via child-abuse hidden services websites, although there is also certainly child-abuse imagery on the normal web. Child abuse sites are another good example of old crimes that now occur in the web-based applications of the Internet.

Gareth Owen and Nick Savage, two computer scientists from the University of Portsmouth in the UK, conducted a highly innovative study to dig into how the Dark Web is used (Owen & Savage, 2015). Their findings highlight the prevalence of child-abuse imagery on the Tor-hosted Dark Web.

For the purposes of their investigation, Owen and Savage volunteered a number of computers into the Tor network. Having put 40 relays at the behest of the Tor network, they then developed a webcrawling bot that automatically read the Tor Hidden Services websites and looked for key words that would allow for the categorization of the hosted darknet websites. Over a six-month period, Owen and Savage categorized the available hidden websites and eventually found that there is a wide distribution of sites on the Tor-hosted Dark Web. As figure 2 shows, the most common type of Tor Dark Web sites relate to drugs, representing something like 15 percent of the sites hosted on the Tor Hidden Services. Those related to child abuse (labeled as 'abuse' in the figure) make up only 2 percent of the sites hosted on the Tor network.


Figure 2. Relative Frequency of Tor's Hidden Services Darknet Sites



Source: Owen and Savage (2015)

On the face of it, Owen and Savage’s initial findings suggest that child abuse is not a big problem on the Dark Web. But the relative frequency of the different types of sites presents only part of the picture. Owen and Savage took the next step in their study and began to track visits to the various Dark Web sites that they had categorized. They found that the distribution of visits to the Dark Web sites is very different from that of hosted sites.

Figure 3. Relative Site Visit Frequency on the Tor Darknet



Source: Owen and Savage (2015)

Once Owen and Savage started tracking the pattern of traffic to the Tor darknet sites, they found that, while abuse sites make up only around 2 percent of the total sites on the Tor Hidden Services, the traffic to this small fraction of sites is disproportionately large. As Figure 3 shows, the pattern of site visits on the Tor-hosted Dark Web sites clusters enormously around the 2 percent of sites dedicated to child abuse. In fact, this small fraction of all the Dark Web sites received roughly 83 percent of the site visits. Certainly, some proportion of this traffic is from automated web crawlers, law enforcement agents

attempting to police the Dark Web and people logging out and then revisiting the site via a different Tor relay. But, even taking into account these refinements of the potential sources of Tor traffic, it is fair to say that the quantitative traffic pattern on the Dark Web tends to cluster around viewing and distributing child abuse imagery online.

As with online illegal marketplaces, the emergence of Dark Web child abuse sites has not created a new form of crime. Instead, the anonymous websites and chat forums of the Dark Web have provided a new and potentially safer location for criminals to conduct their nefarious business. The applications that ride on top of the Internet are the new safe harbours of criminal misdeeds.

Data Breaches and Identity Theft

Falling somewhere in the middle of the continuum of Internet-based crime are acts that combine both the use of the applications on top of the network and the use of the network itself as the primary means of committing the crime in question. Often, crimes in the middle of the range are multistage. These multistage crimes use the infrastructure of the Internet to commit a crime, but then often rely on Dark Web applications to conclude their shadowy business. One such type of crime is data breaches and identity theft.

People have likely been stealing other people's identities from time immemorial. In the decades and centuries before the advent of the Internet, this form of crime involved pretending to be someone else, adopting the person's name, and maybe stealing some sort of credential that might help to legitimize the validity of a stolen identity to society at large. Certainly, this old-fashioned form of identity theft still goes on. Increasingly more common, however, is the theft of someone's digital identity.

Digital identity theft is a two-part process. In the initial step, cybercriminals often use the infrastructure of the Internet to breach a network in order to steal bits of personal information, ranging from Social Insurance Numbers to credit card details, login credentials and email addresses. Unless the crime is targeted and the aim is to steal a particular person's identity, the next phase of the crime involves offloading the stolen information. Selling stolen information is risky, so it usually takes place in the illegal marketplaces of the Dark Web.

The new means of stealing people's personal information take advantage of the increasing digitization of modern life, and the resulting impact can be devastating. More and more, our personal information is stored in large computer-based repositories. Companies and governments alike have moved our information online because there are huge efficiency gains to be had by networking files. It makes finding a person's file and processing requests easier, which streamlines service provision. It also allows users to access their personal records, say, via online banking or online insurance claim websites.

Networked files generate a lot of efficiency gains but, if they are not adequately protected, they also provide a huge boon to those who would steal someone's identity. The problem is that if criminals can hack into a network storing one person's identity, they can often access the confidential information of

a multitude of other people as well. When everything is linked together, a single access point can compromise huge reams of personal data.

A couple of examples drive this point home. The online marketplace eBay was hacked in 2014. The result of a single intrusion that went undetected for months was the compromise of all 145 million eBay user accounts. Luckily, in this case at least, the only data that was exposed was user email accounts (Epstein, 2014). Take another example from 2014. The massive home improvement store Home Depot was hacked via a vendor's credentials. The victims in this instance were not so lucky, as 53 million email addresses and 56 million credit card numbers were taken (Banjo, 2014). In late 2013, the retail giant Target was also hacked via an HVAC (heating, ventilating and air conditioning) vendor's login credentials. Here again, a small breach of the network's defences resulted in some disproportionately large effects. The digital thieves made off with 40 million credit and debit card numbers and tidbits of personal data on up to 70 million people (Smith, 2014).

While the Internet provides the means to break in and steal people's confidential information, the crime of identity theft in the digital age does not end there. The next step is to take the stolen data to an illegal online marketplace, usually hosted on the Dark Web.

On the Dark Web, the price of stolen information varies considerably. The IT security firm McAfee recently published a report on "The Hidden Economy," which provides some insight into the value criminals assign to having access to the details of our digital lives. The authors found that, in the United States, stolen payment cards are worth around 5 to 8 dollars. The price increases as additional biographical information is added to the shopping cart. For the so-called "fullzinfo" version, which includes not just a valid payment card number, but also the "full name, billing address, payment card number, expiration date, PIN number, social security number, mother's maiden name, date of birth, and CVV2" of the card is only 30 dollars per individual profile (McAfee, 2015, p. 5). Clearly, personal data can be had for cheap on the Dark Web.

Stolen personal data is cheap for two reasons. First, the price of stolen credit card details and other bits of personal information like login credentials is variable across the life span of the stolen data. Put simply, credit card numbers expire like produce if not used quickly. Brian Krebs (2014) has shown, for example, that in the wake of the Target hack, a large bundle of credit cards was sold before word of the data breach was reveal for between 26.60 and 44.80 dollars. Only a few months after the Target hack became common knowledge, the price for a similarly sized block of card numbers was somewhere in the range of 8.00 to 28.00 dollars. At the high end, prices had dropped 37.5 percent over this short period, while at the low end prices had fallen by as much as 69.9 percent. The cause of this decline, as Krebs notes, is that credit card details and other login credentials are time delimited. Once people learn that their information has been compromised, they scramble to cancel cards and change passwords rendering the stolen information useless. As the percentage of useful personal data sold in these bundles declines over time, the price of a block of stolen cards goes down quite rapidly.

The low price for bits of stolen information is also a result of an oversupply of sensitive personal data in the marketplace. The online markets for stolen information are markets in the purest sense. As basic economics teaches us, in functioning markets, the interaction of supply and demand determines a commodity's price. The boon of a networked system for criminals is that they can steal a lot of information from a single network breach, while the bane of these cybercriminals is that so can anyone else with the necessary skills and a dose of malicious intent. Barring the emergence of a drastically more secure Internet, the demand for stolen credentials is unlikely to rise faster than its supply; consequently, the price of such a commodity is likely to remain low for the foreseeable future.

Data breaches and online identity theft are a two-step process. The crime of identity theft is committed via the infrastructure of the network, through the use of hacks and social engineering. But, once the private information is in hand, criminals then try to offload their merchandise in the Dark Web forums that exist on top of the network.

Ransomware

Ransomware attacks are also a mixture of the use of the network to transmit malware from one device to another and the use of Dark Web applications to finalize the crime. Ransomware attacks are a bit like a digital version of hostage taking.

In recent years, cryptolocker or cryptowall attacks have emerged as a highly effective form of online criminal extortion. These attacks are based upon the use of ransomware Trojans. Cryptolocker attacks basically involve someone hacking into a person's laptop, desktop, or other device and installing highly sophisticated encryption technologies, which are akin to a really effective safe door. These attacks often start with a simple mistake by an Internet user. Something as seemingly benign as clicking on a link in an email can infect a person's device with a ransomware Trojan. On the face of it, an encrypted device does not seem too harmful, but the trouble is that the only people with the key to the newly locked safe door are the hackers who breached the system. The actual owner of the device is locked out and unable to use or access any of his or her files and applications.

Once the system is locked, the hackers contact the owner of the device via a screen popup and make a ransom demand. The amounts charged are usually fairly small. A typical ransom demand only asks for a modest sum of between 200 and 2,000 dollars, with values in the 300 to 600 dollar range being very common. Late payment penalties can ramp up the amount demanded to as much as 20,000 dollars if the initial ask is not paid by the deadline (Segura, 2013). Payments need to be done via an anonymous currency, such as Bitcoin or Darkcoin, and have to be paid via an online payment website. To avoid detection by any watching law enforcement agencies, the hackers generally host their payment websites on the Dark Web, where everyone is anonymous and communications are heavily encrypted.

There is, however, some honour among thieves. While ransomware puts the owner of the corrupted device in a very vulnerable spot, the hackers almost always honour their word and provide the encryption key to those who make timely online payments (Ducklin, 2015). The motive for providing the encryption keys to those who pay the ransom is most certainly not altruistic. Since no information is

actually stolen during a ransomware attack, the criminals involved need their victim to pay if they are to make any money from their criminal venture. If word was to spread that the criminals behind ransomware attacks were not going to provide their victims with the keys to unlock their systems, people would probably just stop paying and the attack vector would become profitless.

A local Sheriff's department in Dickson, Tennessee, provides an interesting example of how the process usually unfolds. In November, 2014, it was reported that the sheriff's department had run afoul of a form of ransomware known as Cyptowall. The hackers were asking for a paltry, yet oddly specific, sum of 572 dollars. While many of their files were properly backed up, the locked out system did restrict police access to some 72,000 files. It quickly emerged that these files pertained to some ongoing investigations and would otherwise have a negative effect upon the department's ability to serve and protect. The sheriff's department contacted both the FBI and the Tennessee Bureau of Investigation, both of whom were unable to break the highly sophisticated encryption involved. In the end, as the department sheriff Jeff Bledsoe put it, "My first response is we are not going to be held hostage. We are not going to pay a fee to get our records back. But once it was determined which records were involved and that they were crucial to victims of crimes in this county, and to the operations of the sheriff's office and the citizens of this county...I had no choice but to authorize to pay this" (cited in Gadd, 2014). The hackers provided the key and the system was opened.

The use of cryptolocker attacks as a form of online crime has proven highly effective. The IT security firm Norton Symantec notes in their 2015 Internet Security Threat Report that the number of reported cryptolocker attacks grew from 8,274 attacks in 2013 to 373,342 attacks in 2014 (Norton Symantec, 2015). The change amounts to a roughly 45 times (over 4400 percent) increase in one year. Cryptolocker attacks remain the less used cousin of all ransomware attacks, but they are the most effective and they are earning criminals large amounts of money.

Cryptolocker attacks resemble hostage taking in the physical world, but their clever use of the infrastructure of the Internet to deliver malware with the aim of locking a person's computer changes the relationship between the perpetrator and the victim. Cryptolocker attacks also make use of the applications that run on top of the Internet's infrastructure. The money drop website where the victim exchanges some anonymous currency for their device's encryption keys takes place on a Dark Web website.

Distributed Denial of Service Attacks (DDoS)

At the other extreme of the continuum of Internet-based crimes, there are new crimes that rely almost entirely upon the transmission of misdeeds via the network's infrastructure, although there is sometimes a Dark Web component. A quintessential example is the distributed denial of service attack (DDoS).

DDoS attacks capitalize on a weak point inherent to every website, namely, its finite capacity to respond to users' requests to view content. Whenever a person tries to access a website, the webserver goes through a process of noting the request, finding the page that the user wants to see and then displaying

that page to the expectant person. Usually, the process lasts mere seconds. Shortly after the user's click, the desired webpage is displayed.

To varying degrees, all websites are designed to handle a huge volume of user requests every second of every day. The total volume of requests that a website can handle varies depending upon the technical specifications of the server hosting the webpages. Systems with more capacity can, as one would expect, handle more requests per second. But regardless of the size and capacity of a website's servers, there is always a limit to the number of requests that the website can handle at a given point in time. It could be 1,000 requests per second or it could be 100,000, but the upper limit is always there.

DDoS attacks use the power of large numbers of compromised computers to cross this threshold and overwhelm a target by sending in thousands of requests in a short period of time. Once the threshold is passed, the website either slows down significantly, inhibiting a person's ability to access the site, or is knocked completely offline. DDoS attacks get their destructive power by harnessing corrupted computers around the world. The harnessed computers are all infected by malware, which makes their processing power a slave of a remote command and control computer physically located somewhere else in the world and often hosted online on a Dark Web site.

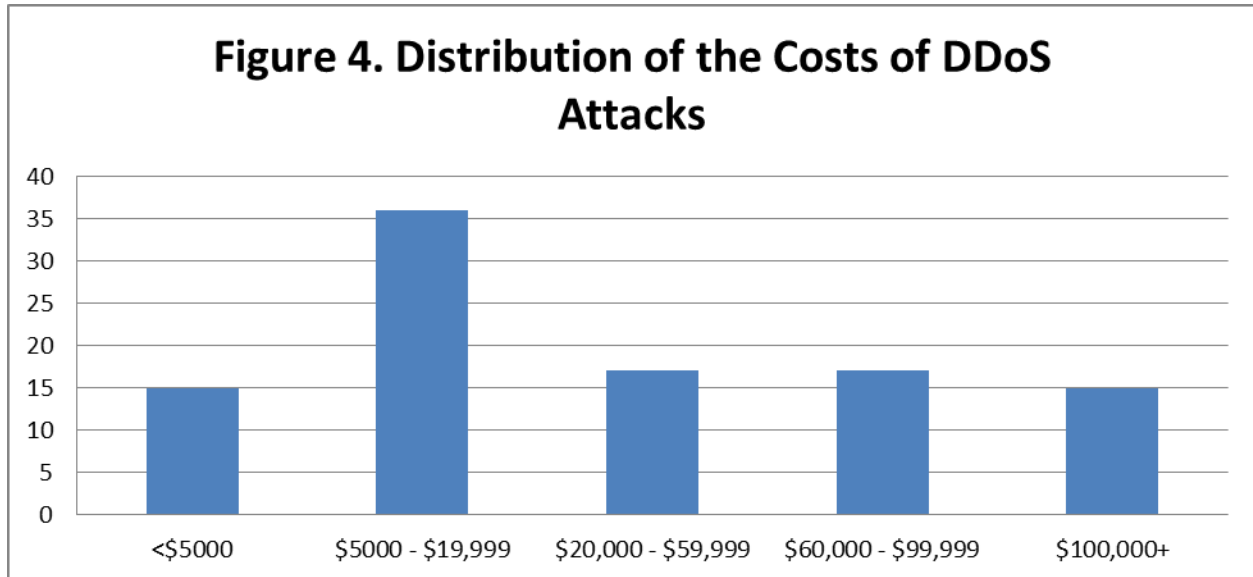
Individual users often are unaware that their computers have been corrupted and harnessed into a botnet. The signs that a user's system has been corrupted are usually very subtle, often manifesting simply as the infected computer running slowly. For the more sophisticated user, evidence that the system is using a lot of its processing power for relatively simple tasks like using Microsoft Word could be taken as a sign that the computer is potentially being used by others for shady purposes.

Botnets can be built from scratch by any malicious actor, but they are more often than not built and then rented out to people with nefarious intent on a piecemeal basis. The services of a botnet are sold in illegal marketplace in the Dark Web. Remarkably, the cost of harnessing the power of a million computers is very affordable. In 2012, Trend Micro noted that an hour-long attack cost 10 dollars, a day-long attack cost 30-70 dollars and a week-long attack cost only around 150 dollars (Trend Micro 2012).

DDoS attacks are a persistent and growing problem. One estimate noted by Verisign, a company involved in, among other things, DDoS mitigation, puts the number of DDoS attacks at around 10,000 separate incidents per day (Verisign, N.D.). The size of DDoS attacks continues to increase year over year, meaning that the ability of DDoS attacks to take down websites is constantly improving (although so is the capacity of websites to handle requests). Verisign again noted, for example, that in Q3 of 2015, the size of DDoS attacks over 10 gigabytes per second went up over 100 percentage points compared to the previous quarter (Verisign 2015, p. 4). This is a bad sign, as larger volume attacks are more powerful and more likely to knock a site offline.

These attacks can be highly disruptive and costly for the targeted businesses. One survey of firms by the firm Incapsula found that the estimated costs of DDoS attacks form a positively skewed distribution (see Figure 4). Around 15 percent of DDoS attacks have an estimated cost to the site owner of less than 5,000 dollars per hour, while a mirroring 15 percent of such attacks cost above an estimated 100,000 dollars

per hour. The average cost of DDoS attacks across these numbers comes in around 40,000 dollars per hour (Matthews, 2014).



Source: Matthews (2014)

DDoS attacks can appear to be a relatively minor annoyance, at least until you view the effect of a shutdown in a wider perspective. Estimates vary, but the duration of DDoS attacks and their ability to take websites down cluster around some clear averages. One study by the Sans Institute, for example, found that the average duration of DDoS attacks in 2013 was 8.7 hours. More importantly, DDoS attacks in 2013 were able to knock websites offline an average of 2.3 hours per event (Pescatore, 2014). Depending upon the site that is taken down, the majority of Internet users might not even notice that a website was knocked offline during a 2.3 hour window, so the effect of DDoS attacks can seem relatively minor.

But the Internet is a very large and active ecosystem, so, even if a majority of people are unaware that a website is down for 2.3 hours, enough users will be affected and the shutdown of a website for even a short period of time can result in very heavy costs. And of course, those costs can be much higher if would-be users have a critical and urgent need to interact with the website. Multiplying together the average number of attacks per day (10,000 attacks), the average downtime due to attacks (2.3 hours) and the average cost of the DDoS attack (40,000 dollars) can give a sense of the economy-wide costs of this sort of online crime. While you need to take a multiplication of averages like this with caution, crunching the numbers indicates that the rough, average cost of DDoS attacks come in around 920,000,000.00 dollars per day. While these numbers are probably far too high, they do illustrate the fact that, in a network of Internet's size and scope, even small costs can accumulate rapidly.

DDoS attacks are like a digital version of a picket line, strike or boycott. Their aim is to disrupt a company or government's ability to provide a service. While the botnet command and control nodes that manage the computers that are laced together to launch a DDoS attack are often hosted on the Dark Web, the crime itself is committed via the infrastructure of the Internet.

Internet-Based Crime and Two Policy Response Examples

Internet-based crime ranges along a continuum from crimes that have shifted into the applications that ride on top of the Internet to crimes that manipulate the infrastructure of the Internet to transmit misdeeds around the world at the speed of light. Understanding that online crimes rely upon the technology of the Internet in different ways is important for policy makers who aim to devise crime mitigation strategies. A policy aimed to address one type of crime is quite likely to spill over into other areas because online crime is based upon, broadly speaking, just two uses of the technology: the commission of crime in the applications on top of the network and crime that is committed via the infrastructure of the network.

Combating crime in the digital age is no longer the sole purview of governments. Increasingly, governments are teaming up with private sector actors, including large technology companies like Microsoft and IT security companies like Norton Symantec. Beyond governments, individuals and Internet companies like Internet service providers (ISPs) are also involved in preventing cybercrime. It is a bit of an all-hands-on-deck situation, with each party doing what it can (or, rather, what it has an incentive to do) to prevent or mitigate Internet-based crime in all its manifestations.

In such a complex ecosystem, different kinds of policies can be used to try to counter Internet-based crime. Any given policy might have a large effect on a particular type of Internet-based crime, but a small effect on another type of criminal behaviour. It is also possible that the 'sign' of the causal relationship might reverse itself as we move from one type of crime to another. Thus, for example, a policy that aims to help reduce crime committed in the applications that ride on top of the Internet might actually make it easier to commit crime via the network by diminishing the Internet's level of technical security. Government back doors in encrypted messaging applications would be one such example, as such an effort would create a number of technical challenges that could make breaching systems and stealing information easier (Abelson et al., 2015).

Table 1 presents a summary of how effective two particular crime-mitigation policies can be in countering various forms of Internet-based crime. To the right of each policy, the table lays out the sign (positive or negative) and the probable magnitude (large, medium or small) that the policy is likely to produce for each of the five examples of Internet-based crime discussed in the previous sections. In the next two subsections, I go through the likely effects that each crime-fighting strategy is likely to have on the various types of crime analyzed above.

Table 1. Likely impact of two crime-mitigation strategies on five types of Internet-based crimes

	Illegal Marketplaces	Child Abuse Sites	Data Breaches/ Identity Theft	Ransomware	DDoS Attacks
Dark Web Indexing	Large	Large	Medium	Small	Medium
ISP Botnet Mitigation	Negligible/Not Applicable	Negligible/Not Applicable	Small	Medium	Large

Dark Web Indexing

All five examples discussed in the previous section involve the Dark Web in some capacity. Illegal markets and child-abuse sites involve the Dark Web most directly, as the occurrence of crime online has shifted to the Dark Web as a platform for criminal activity. Data breaches and identity theft also involve illegal Dark Web marketplaces in the sale phase of the criminal undertaking. Even the strict online crimes (ransomware and DDoS attacks) involve the Dark Web. Ransomware attacks set up Dark Web exchange sites so that people who are locked out of their systems can pay the criminals. The owners and operators of the botnets used in DDoS attacks also rely on the Dark Web to host their command and control nodes, so as to provide some distance between their personal devices and the command architecture of illegal botnets.

Law enforcement agencies have responded to the growing challenges of the Dark Web by trying to identify and index the sites that are out there. Often, law enforcement agencies team up with private sector actors to try to figure out what the contours of the Dark Web are. The international crime-fighting agency Interpol, for example, has teamed up with the IT security firm Kaspersky Labs to index Dark Web sites. Interpol has indexed around 20,000 sites, and Kaspersky has indexed another 5,000. These are but a small fraction of the estimated 60,000 to 80,000 sites that exist on just the Tor Hidden services, not to mention the other similar hosts of Dark Web material.

Figuring out the contours of the Dark Web has effects across the spectrum of Internet-based crime, although the magnitude varies, for a number of reasons. To put it simply, knowing where the crime is taking place is an essential prerequisite for effective law enforcement. To use a simple analogy, if the police are not sure where neighbourhood drug dealers congregate in the physical world, their efforts to counter drug traffic will be highly inefficient at best, and wide of the mark at worst.

The same basic logic applies on the Dark Web. Effectively policing illegal marketplaces and child abuse sites requires, first and foremost, that the police know the location of the sites that are being used as forums for criminal activity. From there, the policy can take active measures (from exploiting the anonymity of the system to exploiting vulnerabilities in the technology) to try to identify and track down the criminals involved.

The policy is particularly effective for policing illegal markets and child abuse sites because these websites are often subject to a relatively low churn rate when compared to other Dark Web sites (Owen and Savage, 2015). Since they typically stay up for a long period of time, knowing where these sites are located can be a significant boon for law enforcement agents. To the extent that identity theft relies on stable Dark Web marketplaces to offload stolen credentials, Dark Web indexing can also help to deter subsequent identity misuses and abuses. Botnet command and control nodes need to be somewhat stable in order for the infected computers to report for orders from their master. Pinpointing botnet command and control nodes can facilitate their takedown by law enforcement agencies. Across this wide range of criminal undertakings, Dark Web indexing can be quite effective.

The one outlier is likely ransomware attacks, which are probably hardly affected at all by Dark Web indexing. The sites used for the ransom money drop are, potentially at least, highly transient. A single site can be used for multiple money drops, but they are not particularly costly or time-consuming to set up. Ransomware thieves can simply move on to a new site after every transaction. The perpetrators of the attacks provide the victim with the website address of the drop. They neither need nor necessarily want a stable locale for their transactions. Thus, indexing Dark Web sites is likely to only have a very small mitigating effect on the commission of ransomware attacks.

In short, Dark Web indexing can be a very useful tool in the law enforcement official's toolkit, but its effect across the various forms of Internet-based crime is variable. It tends to be highly effective as a tool to counter illegal marketplaces and child abuse rings that relies upon the stable use of Dark Web applications. Dark Web indexing will likely only have a medium-sized impact on the mitigation of identity theft and DDoS attacks, as these forms of criminal behaviour do rely upon Dark Web websites but are able to launch new sites if law enforcement tracks them down. And indexing is unlikely to have any serious impact on countering ransomware attacks.

ISP Botnet Mitigation Strategies

Malicious botnets are the *demons* behind the DDoS attacks that knock banking, government and newspaper websites offline. They are also often used to launch massive email spam campaigns that can spread malware, result in fraud and waste everyone's time. As Dave DeWalt, the former CEO of McAfee, put it, botnets are "the engine that drive everything" malicious (cited in Rowe et al., 2011, p. 2).

Individuals and organizations alike rely upon Internet Service Providers (ISPs) to transmit their Internet traffic across the wider Internet. Many of these networks are privately owned, and some are run by state corporations. While there are billions of Internet users and even more devices online, there are

relatively few ISPs. By dint of their relatively small numbers, ISPs are chokepoints for Internet traffic flows and can be useful actors in the fight against Internet-based crime.

Just as there are good citizens and bad citizens in the offline world, there are also good netizens and bad netizens online. ISPs are no different. While some ISPs work to provide security to their users and to prevent the use of their infrastructure for botnet activity, other ISPs let their networks be used for criminal purposes. One study that looked at the origins of spam traffic as a means of backtracking botnet computers to their hosting ISP network illustrates this point well. The authors found that there are indeed a few digital laggards that reduce the level of cybersecurity for everyone (Van Eeten et.al. 2010). In particular, the worst 10 ISPs in term of hosting botnet computers actually account for something on the order of 30 percent of all unique IP addresses sending spam. The worst 50 ISPs account for over half the total spam-sending IP addresses. With a margin for error, these results suggest that a small number of networks provide Internet access to most of the botnets in cyberspace. These numbers are startling in “light of the fact that there are ... anywhere between 4,000 -100,000 ISPs” worldwide (Ibid., p. 9).

Malicious botnets are clearly clustered on a few troublesome ISP networks. As a result, ISPs can be leveraged to block botnet traffic of dubious provenance and thus reduce the severity of spam and DDoS campaigns. If the incentives can be correctly aligned, ISPs can engage in effective botnet remediation strategies. For instance, by looking for irregular traffic patterns, ISPs can pinpoint infected computers and route their traffic into sinkholes so that no users get exposed to malicious activity of any form. They can also contact the account holders of infected devices on their network to inform them that their system has likely become a part of a botnet and suggest ways to remedy the problem. Melissa Hathaway, a former Whitehouse cybersecurity expert, said it best when she noted that the effect of this sort of ISP-driven botnet mitigation strategy aims to “‘drain the swamp’ of malicious cyber activity and tilt the playing field [back] in our favour” (cited in Singer & Friedman, 2014, p. 179).

As indicated in table 1, ISP-driven security policy will have different levels of effectiveness across the various types of Internet-based crime. ISP botnet mitigation strategies can be highly effective at reducing DDoS attacks, as such attacks hinge upon the numbers that botnets provide. The strategy of mitigating botnet activity can also be potentially effective at limiting, to a certain extent, ransomware attacks. The emails that contain the malicious links that unfurl into an encrypted device and a demand for ransom can be delivered as targeted phishing emails, but they can also be sent via more mass-based template emails (similar in form to spam). To the extent that spam contains online trickery designed to pilfer away people’s private information, botnet mitigation strategies can also help here. Yet ISPs’ efforts to mitigate botnet activity will have nearly no effect on the occurrence of crime in the applications of the Dark Web. Dark Web illegal marketplaces do not rely on botnets and so will be immune to the crime containment efforts of ISPs. The same applies to child abuse sites.

Conclusions

In this chapter, I have developed the notion that effective cybersecurity policy requires a clear understanding of the type of Internet-based crime with which law-enforcement officials are trying to contend. I have argued further that Internet-based crime falls along a continuum ranging from crimes that have shifted into the applications that ride on top of the Internet to criminal misdeeds that are carried out via the infrastructure of the network.

To illustrate the usefulness of conceptualizing cybercrime along a continuum, I have first discussed how illegal marketplaces and websites hosting child abuse imagery use the applications and forums of the Dark Web as the new arena of illegal activity, effectively replacing the old gathering places of the offline world. This sort of crime relies upon the applications that ride on top of the infrastructure of the Internet. I have then discussed how data breaches and ransomware attacks involve separate stages of criminal wrongdoing. In both cases, the first step involves using the infrastructure of the Internet to breach a victim's system, either to steal information or hold the data ransom, respectively. The second stage involves monetizing the criminal misdeed, which usually involves a transaction in a Dark Web application such as an illegal market or a specially setup ransom drop site. Lastly, I pointed to how DDoS attacks tend to only involve the use of the infrastructure of the Internet for criminal purposes, but which do get rented out via illegal markets and that do hinge upon the operation of a botnet command and control node hosted somewhere upon the Dark Web.

The third section of the chapter discussed how cybersecurity policies designed to help police cyberspace have different levels of effectiveness in mitigating the various types of crime. Indexing Dark Web sites has a salutary effect on mitigating all the types of crime considered here, but it is likely to be most effective in mitigating crimes that use Dark Web applications as forums for illegal activity (e.g., illegal markets and child abuse sites). It will likely have a more modest effect on efforts to counter identity theft and data breaches and DDoS attacks. The impact of Dark Web indexing on efforts to mitigate ransomware is again likely to be positive, but also likely quite small. In other words, the effectiveness of Dark Web indexing is ranged from the largest mitigating impact on crimes that occur in digital applications to the smallest impact on crimes that are committed via the infrastructure of the Internet.

The second policy response that I discussed was ISP-led botnet mitigation strategies involving traffic redirection and notification of those individuals whose devices have been compromised and who have become part of a malicious botnet. Combatting botnets on the Internet can greatly contribute to the reduction of DDoS attacks. To the extent that ransomware and identity theft result from phishing (targeted emails that try to get people to click on links or attachments) or spam email campaigns, blocking botnet traffic at an ISP level can again be highly effective. On the other hand, botnet mitigation strategies probably have no effect on the occurrence of crime in online applications. In other words, this crime mitigation strategy is most effective against criminal misdeeds that occur via the infrastructure of the Internet and has almost no effect on crime that had shifted to the applications that run on top of the Internet.

The approach taken here suggests several avenues for future scholarship on cybercrime. Other typologies for computer-based crime certainly exist. Wall (2001), for example, proposes a typology that breaks cybercrime into four categories: cyber-trespass, cyber-deception/theft, cyber-porn and obscenity, and cyber-violence (for an application, see Holt & Bossler, 2014). Wall's typology divides crime according to the nature of the criminal acts. As befits a good typology, Wall's framework is exhaustive, but it treats each type of crime as discrete from every other type.

The advantage of the current framework is that it divides the commission of Internet-based crime according to how the crime itself relies upon the Internet as a technology. The virtue of this approach is that it allows for a clear articulation of the similarities between otherwise dissimilar forms of crime. For example, data breaches and child abuse sites are similar in the sense that they both rely upon Dark Web applications. The degree of dependence is different, of course, but the reliance is there. It also allows for a clear articulation of the dissimilarities between different crimes. DDoS attacks, for example, are similar to online illegal markets in that there is a Dark Web connection, but they are different in that the former is committed solely via the infrastructure of the Internet while the latter occurs on top of the network. These similarities and differences are useful for policy makers to know.

The framework provides a number of avenues for future research. First, the reliance to varying degrees of all types of crime on the Dark Web suggests that more needs to be done to understand both the motives driving people to use anonymizing technologies such as Tor and how certain Dark Web sites gain in popularity. For example, why did Silk Road emerge as the dominant illegal marketplace when there were so many similar alternatives? Knowing why clusters of activity form on the Dark Web can aid law enforcement in disrupting these processes, mitigating the occurrence of cybercrime in web-based applications.

Second, more extensively mapping out the ways in which crime-mitigation strategies affect different crimes along the continuum would provide tremendous utility for law enforcement agencies. It would provide useful insight into how a single strategy may span different areas. It could also reveal how a policy that is developed in a myopic environment to counter one type of crime might have counterproductive effects on other areas in clearly predictable ways, as back doors in encryption would likely have on web security.

Overall, Internet-based crime is best understood as varying along a continuum based upon how the criminals actually use the network. Some crime rides on top of the Internet, making use of the system's applications and forums. Other crimes are committed via the infrastructure of the Internet itself. This distinction helps clarify how crime interacts with efforts to mitigate it and how policymakers can hopefully move towards a more secure cyberspace.

References

- Abelson, H. Anderson, R. Bellovin, S.M. Benaloh, J. Blaze, M. Diffie, W. Gilmore, J. Green, M. Landau, S. Neumann, P.G. Rivest, R.L. Schiller, J.I. Schneier, B. Specter, M., & Weitzner, D.J. (2015). Keys under doormats: mandating insecurity by requiring government access to all data and communications. *Computer Science and Artificial Intelligence Laboratory Technical Report*. Retrieved from <https://dspace.mit.edu/bitstream/handle/1721.1/97690/MIT-CSAIL-TR-2015-026.pdf>
- Ax, J. (2015). A Silk Road drug-dealer turned government-witness was just sentenced to over 2 years in prison. *Business Insider*. Retrieved from <http://www.businessinsider.com/r-silk-road-drug-dealer-turned-government-witness-gets-2-12-years-in-prison-2015-7>
- Banjo, S. (2014). Home Depot hackers exposed 53 million email addresses. *Wall Street Journal*. Retrieved from <http://www.wsj.com/articles/home-depot-hackers-used-password-stolen-from-vendor-1415309282>
- Bartlett, J. (2014). *The Dark Net: Inside the Digital Underworld*. London: William Heinemann.
- Christin, N. (2012). Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online mMarketplace. Working paper. Retrieved from <http://arxiv.org/pdf/1207.7139.pdf>.
- CIGI/Ipsos. (2014). Global Survey of Internet Security and Trust. Retrieved from <https://www.cigionline.org/internet-survey>

- Dean, D., Digrande, S., Field, D., Lundmark, A., O'Day, J., Pineda, J., & Zwillenberg, P. (2012). The Internet Economy in the G-20: The 4.2 Trillion Dollar Growth Opportunity. *Boston Consulting Group*. Retrieved from <https://www.bcg.com/documents/file100409.pdf>
- Ducklin, P. (2015). Ransomware – Should you Pay?. *Naked Security*. Retrieved from <https://nakedsecurity.sophos.com/2015/03/19/ransomware-should-you-pay/>
- Epstein, Z. (2014). eBay thought user data was safe, but 145 million accounts were compromised in massive hack. *BGR*. Retrieved from <http://bgr.com/2014/05/27/ebay-hack-145-million-accounts-compromised/>
- Gadd, C. (2014). Dickson Sheriff's Office pays ransom to cyber criminals. *The Tennessean*. Retrieved from <http://www.tennessean.com/story/news/local/dickson/2014/11/11/dickson-sheriffs-office-pays-ransom-cyber-criminals/18868325/>
- Greenberg, A. (2015). A heroin dealer tells the Silk Road jury what it was like to sell drugs online. *Wired*. Retrieved from <http://www.wired.com/2015/01/silk-road-heroin-dealer-testifies/>
- Hern, A. (2015). Stolen credit card details available for £1 each online. *The Guardian*. Retrieved from <http://www.theguardian.com/technology/2015/oct/30/stolen-credit-card-details-available-1-pound-each-online>
- Holt, T.J. and Bossler, A.M. (2014). An assessment of the current state of cybercrime scholarship. *Deviant Behaviour*, 35(1), 20-40.
- Jardine, E. (2015a). The Dark Web dilemma: Tor, anonymity and online policing. *Global Commission on Internet Governance Paper Series No. 21*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2667711
- Jardine, E. (2015b). Global cyberspace is safer than you think: real trends in cybercrime. *Global Commission on Internet Governance Paper Series No. 16*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2634590
- Jardine, E. (N.D.) Tor, what is it good for? political repression and the use of online anonymity-granting technologies. *New Media & Society* (Under review).
- Krebs, B. (2014). Fire sale on cards stolen in Target breach. *Krebs On Security*. Retrieved from <http://krebsonsecurity.com/2014/02/fire-sale-on-cards-stolen-in-target-breach/>
- Matthews, T. (2014). Incapsula Survey: what DDoS attacks really cost businesses. Retrieved from <http://lp.incapsula.com/rs/incapsulainc/images/eBook%20-%20DDoS%20Impact%20Survey.pdf>
- McAfee. 2015. The Hidden Economy: The Market Place For Stolen Digital Information. Retrieved from <http://www.mcafee.com/us/resources/reports/rp-hidden-data-economy.pdf>

- Norton Symantec. (2015). Internet Security Threat Report. Retrieved from http://www.symantec.com/security_response/publications/threatreport.jsp
- Owen, G. and Savage, N. (2015). The Tor Dark Net. *Global Commission on Internet Governance Paper Series* No. 20. Retrieved from <https://ourinternet.org/publication/the-tor-dark-net/>
- Pélissié du Rausas, M., Manyika, J., Hazan, E., Bughin, J., Chui, M., and Said, R., (2011). Internet Matters: The Net's Sweeping Impact on Growth, Jobs and Prosperity. *McKinsey & Company*. Retrieved from http://www.mckinsey.com/insights/high_tech_telecoms_internet/internet_matters
- Pescatore, J. (2014). DDoS attacks advancing and enduring: a SANS survey. *SANS Analysis Survey*. Retrieved from <https://www.sans.org/reading-room/whitepapers/analyst/ddos-attacks-advancing-enduring-survey-34700>
- Rowe, B., Wood, D., Reeves, D., and Braun, F. (2011). The role of Internet service providers in cybersecurity. *Research Brief*. Institute for Homeland Security Solutions. Retrieved from <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.473.2323&rep=rep1&type=pdf>
- Segura, J. (2013). CryptoLocker ups the ante, demands \$2,000 for overdue ransom. *MalwareBytes Unpacked*. Retrieved from <https://blog.malwarebytes.org/cyber-crime/2013/11/cryptolocker-ups-the-ante-demands-2000-for-overdue-ransom/>
- Singer, P.W., and Friedman, A. (2014). *Cybersecurity and Cyberwar: What Everyone Needs to Know*. Oxford: Oxford University Press.
- Smith, C. (2014). Expert who first revealed massive Target hack tells us how it happened. *BGR*. Retrieved from <http://bgr.com/2014/01/16/how-was-target-hacked/>
- Trend Micro. (2012). Russian underground 101. Retrieved from <http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-russian-underground-101.pdf>
- Van Eeten, M., Bauer, J., Asghari, H., Tabatabaie, S., & Rand, D. (2010). The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data. Retrieved from http://www.econinfosec.org/archive/weis2010/papers/session4/weis2010_vaneeten.pdf
- Verisign. (2015). Verisign distributed denial of service trends report. 2(3). Retrieved from https://www.verisign.com/en_US/security-services/ddos-protection/cyber-security-resources/index.xhtml
- Verisign. (N.D.). Fact sheet: verisign ddos protection services. Retrieved from <https://www.verisign.com/assets/datasheet-ddos-overview.pdf>
- Wall, D. (2001). Cybercrimes and the Internet. In David S. Wall (Ed.), *Crime and The Internet* (pp.1-17). New York: Routledge.

