CIGI

CHATHAM HOUSE
The Royal Institute of International Affairs

Global Commission on Internet Governance

ourinternet.org

PAPER SERIES: NO. 16 — JULY 2015

Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime

Eric Jardine

GLOBAL CYBERSPACE IS SAFER THAN YOU THINK: REAL TRENDS IN CYBERCRIME

Eric Jardine

CHATHAM HOUSE
The Royal Institute of International Affairs

Published by the Centre for International Governance Innovation and Chatham House.

67 Erb Street West
Waterloo, Ontario N2L 6C2
Canada
tel +1 519 885 2444 fax +1 519 885 5450
www.cigionline.org

10 St James's Square
London, England SW1Y 4LE
United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org

TABLE OF CONTENTS

ABOUT THE GLOBAL COMMISSION ON INTERNET GOVERNANCE

The Global Commission on Internet Governance was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. The two-year project conducts and supports independent research on Internet-related dimensions of global public policy, culminating in an official commission report that will articulate concrete policy recommendations for the future of Internet governance. These recommendations will address concerns about the stability, interoperability, security and resilience of the Internet ecosystem.

Launched by two independent global think tanks, the Centre for International Governance Innovation (CIGI) and Chatham House, the Global Commission on Internet Governance will help educate the wider public on the most effective ways to promote Internet access, while simultaneously championing the principles of freedom of expression and the free flow of ideas over the Internet.

The Global Commission on Internet Governance will focus on four key themes:

- enhancing governance legitimacy — including regulatory approaches and standards;
- stimulating economic innovation and growth — including critical Internet resources, infrastructure and competition policy;
- ensuring human rights online — including establishing the principle of technological neutrality for human rights, privacy and free expression; and
- avoiding systemic risk — including establishing norms regarding state conduct, cybercrime cooperation and non-proliferation, confidence-building measures and disarmament issues.

The goal of the Global Commission on Internet Governance is two-fold. First, it will encourage globally inclusive public discussions on the future of Internet governance. Second, through its comprehensive policy-oriented report, and the subsequent promotion of this final report, the Global Commission on Internet Governance will communicate its findings with senior stakeholders at key Internet governance events.

www.ourinternet.org

ABOUT THE AUTHOR

Eric Jardine joined CIGI as a research fellow in May 2014 in the Global Security & Politics Program. He contributes to CIGI's work on Internet governance, including the CIGI–Chatham House-sponsored Global Commission on Internet Governance. His current research focuses on cyber security, cyber terrorism, cybercrime and cyber protest. He holds a Ph.D. in international relations from the Norman Paterson School of International Affairs at Carleton University.

In politics, what begins in fear usually ends in folly.

— Samuel Taylor Coleridge

EXECUTIVE SUMMARY

What are the real trends in cybercrime? Recent media coverage has been rife with stories of large-scale data breaches, hacks and online financial crime. Information technology (IT) security firms such as Norton Symantec and Kaspersky Labs publish yearly reports that generally show the security of cyberspace to be poor and often getting worse. This paper argues that the level of security in cyberspace is actually far better than the picture described by media accounts and IT security reports. Currently, numbers on the occurrence of cybercrime are almost always depicted in either absolute (1,000 attacks per year) or as year-over-year percentage change terms (50 percent more attacks in 2014 than in 2013). To get an accurate picture of the security of cyberspace, cybercrime statistics need to be expressed as a proportion of the growing size of the Internet (similar to the routine practice of expressing crime as a proportion of a population, i.e., 15 murders per 1,000 people per year). To substantiate this argument, data was collected on the size of the Internet, with a focus on users, points of interaction and volume of online activity. Data was then collected on the vectors of cyber attack, the occurrence of cyber attacks and the cost of cybercrime. Normalizing these crime statistics around various measures of the growing size of cyberspace, a clear picture emerges: the absolute numbers always paint a worse scenario of the security of cyberspace than the normalized numbers. In particular, the absolute numbers tend to lead to one of three misrepresentations: first, the absolute numbers say things are getting worse, while the normalized numbers show that the situation is improving; second, both numbers show that things are improving, but the normalized numbers show that things are getting better at a faster rate; and third, both numbers say that things are getting worse, but the normalized numbers indicate that the situation is deteriorating more slowly than the absolute numbers. Overall, global cyberspace is actually far safer than commonly thought.

INTRODUCTION

Recent media coverage has been chock full of high-profile accounts of cybercrime. Hacks, data breaches, destruction of property and the theft of personal information seems to be rampant. In February 2014, eBay's online system was breached after some of its employees' credentials were stolen, leading to the compromise of some 145 million account holders (Finkle, Chatterjee and Maan 2014). In July, the American bank JPMorgan Chase was hacked, with online bandits making off with account information on approximately 76 million households and some eight million small businesses (Silver-Greenberg, Goldstein and Perlroth 2014). In November, Sony Pictures was subject to a sophisticated cyber attack, causing massive physical damage to its computer systems and exposing sensitive emails regarding pay disparities and personal relationships. In December 2014, Sony estimated that the remediation and investigation costs of the hack could enter into the $100 million[1] range (Richwine 2014). What is more, these are just a few of the publicly known breaches.

As the Internet comes to underwrite more and more of our daily life, the vectors of attack for cybercriminals, hackers and state officials multiply, the total number of cyber attacks grows year over year and the potential damage from cyber attacks increases. Governments, corporations and individuals have prudently responded to these trends by stepping up their cyber defences. Shortly after the Sony Pictures hacks, for example, the United States and the United Kingdom announced a series of "cyber war games" to prepare their government agencies for the potential of broad-based cyber attacks on critical infrastructure, including the banking and financial sector (BBC News 2015). Over 60 percent of businesses' representatives surveyed in a recent Gandalf Group C-Suite study have responded to the perception of a deteriorating cyber security environment by increasing their IT security budgets (Gandalf Group 2014). Likewise, a recent CIGI-IPSOS poll surveying over 23,000 respondents in 24 countries found that 64 percent of respondents were more worried about their online privacy compared to one year ago and 78 percent of respondents were concerned about criminal hackers stealing their banking information. An additional 77 percent of respondents were concerned that online criminals would steal their private messages and photos. Indicating the behavioural changes that people have undertaken in response to perceptions of the poor security of cyberspace, the survey also found that compared to one year ago, some 43 percent of respondents now avoid certain Internet sites and web applications, about 39 percent change their passwords regularly and roughly 10 percent actually use the Internet less often (CIGI-IPSOS 2014).

Clearly, the proliferation of cybercrime and the media's coverage of high-profile hacks have generated a severely negative perception of the security of cyberspace and caused governments, business and individual citizens to take additional steps to protect themselves online. The problem is that the existing picture of the security of cyberspace is misleading. Currently, statistics on cybercrime are, as far as I am aware, always expressed in either absolute (1,000 attacks per year) or year-over-year (2013 had 46 percent

1 All currency is in US dollars.

more cyber attacks than 2012) terms.[2] The difficulty with this expression of the numbers is that it gives an inaccurate picture of the actual trends in cybercrime over time, and thus a false impression of the actual security of cyberspace. To state the obvious (but perhaps not well understood), the occurrence of cybercrime is inevitably related to the size of the Internet. Since cyberspace is, in a number of ways, expanding at an exponential rate, it is reasonable to expect that the *absolute* number of cyber attacks will also increase simply because the Internet ecosystem is getting bigger and not necessarily because the situation is growing worse. These observations raise two questions: What is the actual trend in cyber security? And is cyberspace becoming less safe, safer or staying roughly the same over time?

In order to provide an accurate picture of the security of cyberspace, all indicators of cybercrime need to be normalized around data that captures the growing size of the Internet.[3] An example to help clarify the importance of normalizing (or, essentially, expressing numbers as a proportion of a population) data on cybercrime around the size of the Internet is as follows: Imagine there is a town of 1,000 people with 100 violent crimes a year. Now imagine that there is a city with 100,000 people with 1,000 violent crimes per year. When normalizing the crime statistics for these two hypothetical population centres, it is found that the town has a violent crime rate of 0.1, while the city has a violent crime rate of 0.01. In other words, even though the city has as many violent crimes as the entire population of the town, a person's chance of being subject to a violent crime in the city is only 1 in 100, while the chance of being the victim of a violent crime in the town is 1 in 10.

In the case of the global Internet, the occurrence of cybercrime can only be meaningfully normalized around figures that capture the full width and breadth of cyberspace. Cyber attacks in one country can originate in any other country on the planet that has an Internet connection. Normalizing crime statistics around national-level data, therefore, gives a partial and highly skewed glimpse at real trends in the occurrence and cost of cybercrime.

Taking data on the size of the Internet and normalizing various cybercrime indicators around these figures from 2008 to the end of 2014, the security of cyberspace is better than one would think from looking at just the absolute numbers often presented in the media and in IT security reports. Over 30 comparisons of the absolute (1,000 attacks) and normalized (0.15 attacks per 1,000 Internet users) numbers bear out this claim.

When the normalized indicators of cybercrime are compared to the absolute numbers that are usually used to discuss the level of security in cyberspace, one of three misrepresentations occurs:

- the absolute numbers indicate the situation is getting worse when the normalized numbers say it is getting better (as in the case of new vulnerabilities, zero-day vulnerabilities, browser vulnerabilities, mobile vulnerabilities, post-breach response costs and notification costs);
- both the absolute and the normalized numbers say the situation is worsening, but the absolute numbers say it is growing worse at a faster rate than the normalized numbers (as in the case of detection and escalation costs, when the full sample is considered); or
- both the absolute and the normalized numbers say the situation is improving, but the absolute numbers indicate a slower rate of improvement than the normalized numbers (as in the case of malicious web domains, botnets, web-based attacks since 2012, average per capita data breach costs, organizational costs due to data breaches, detection and escalation costs from 2010 to 2013 or lost business costs).

In short, when the number of cyber attack vectors, the number of cyber attacks and the amount of damage caused by cybercrime are expressed as a proportion of the size of the Internet, each of the normalized numbers point to the idea that the security of cyberspace is better than is suggested by the un-normalized or absolute numbers. As a result, the security of cyberspace is likely better than is commonly perceived by the general public, private companies and state officials.

A realistic understanding of the level of security in cyberspace is important because an unnecessarily negative image of the situation can lead to radical policy responses that could easily produce more harm than good. If online crime is rampant, then restricting online activity might be warranted, likely to the ultimate detriment of cultural expression, commerce and innovation. If, on the other hand, cyberspace security is relatively good, then current policies could be sufficient and things can go on more or less as they do now. In any case, a more realistic impression of the security of cyberspace provides a better foundation for cyber security policy.

2 The two exceptions involve spam and phishing emails, often expressed as a percentage of all emails sent. There is no clear rationale given for why cybercrime statistics are expressed in absolute or year-over-year terms. One potential reason is that, as shown in this paper, the numbers tend to be more severe and point to a worse situation. Since most collectors of cybercrime data are private, for-profit companies, a cynic could conclude that the companies present data in a specific way to help them sell product. I have no proof at all of this interpretation. It is merely one potential explanation.

3 In this paper, the terms Internet and cyberspace are used synonymously. The Internet usually refers to the physical structure of the network, while cyberspace is the larger, over-the-top portion of the Web involving things such as apps. Both terms herein mean cyberspace and both are used in the paper to mean the same thing in the interest of readability.

The paper first discusses how to conceptualize the size of cyberspace and details the data that is used herein to measure this concept. It then provides a three-part framework for thinking about the security of cyberspace and details the measures used to operationalize each part of the framework. The next three sections examine the normalized trends in each of these areas and compares them to the trends in the absolute numbers. The paper concludes with policy recommendations based on the finding that cyberspace security is better than what is indicated when looking at only the absolute numbers and is actually, in many cases, getting better rather than worse.[4]

THE SIZE OF CYBERSPACE

The cyberspace ecosystem is built upon the physical infrastructure of the Internet and is basically composed of users, points of online interaction (websites, for instance) and the volume of activity that occurs online. The online ecosystem gets larger as the number of users, points of interaction and volume of activity increases. This section lays out a three-part framework for understanding the scope, size, width and breadth of cyberspace. Cyberspace is essentially an amalgamation of the number of users (people and devices, etc.), the number of points of interaction (websites and domains, etc.) and the activity linking these broad categories (data flows and commerce, etc.).[5]

The basic point is that the ecosystem of cyberspace is big and getting a lot bigger at a fairly rapid pace. This growth is akin to the growth of a population in a city or country, in the sense that a fixed amount of crime and a growing population will result in a lower crime rate or a better chance that one will not be subject to a crime.

As detailed below, data was collected from a variety of sources on the following variables for the concept of Internet users:

- the number of Internet users;
- the number of email users;
- the number of active mobile broadband subscriptions; and
- the number of smartphones sold to end-users.

The following data was collected on the concept of points of online interaction:

- the number of domains; and
- the number of websites.

And on the volume of online activity:

- the volume of total data flows;
- the volume of mobile data flows;
- the annual number of Google searches; and
- the Internet's contribution to GDP.

Table 1 provides some basic summary statistics for the data capturing the size of cyberspace.

Internet Users

The number of Internet users is a good measure of the size of cyberspace because it shows the actual number of people that are a part of the "network of networks." In this sense, it is akin to the number of people in a city or country. It is also a good proxy for the number of devices online, although this number surpassed that of humans on the network around 2008 (Evans 2011). Data on the number of Internet users from 2008 to the end of 2014 was taken from the website Internet Live Stats, which provides real-time statistics on various indicators of the size of the Internet (Internet Live Stats 2015a).

Email is one of the most basic uses of the Internet. The number of email users online is a good measure of the size of the active population base of the online ecosystem because it captures not just the number of people who have web access (as done via Internet users statistics), but also the number of people who actually use the Internet as a part of their daily lives. Email users, therefore, are an active subset of all Internet users. In 2014, for example, there were 421,249,355 more Internet users than email users for that year. Data on email users from 2008 to 2012 was taken from a data aggregation blog called Royal Pingdom, which is operated by the website monitoring company Pingdom (Royal Pingdom 2009; 2010; 2011; 2012; 2013). Data for email users for 2013 and 2014 were taken from a Radicati Group (2013) study of the email market.

Increasingly, people access the Internet via a mobile platform rather than a traditional desktop computer. In January 2014, mobile usage surpassed desktop usage in the United States for the first time (O'Toole 2014). The trend is even more pronounced in the developing world, where Internet access has expanded primarily by skipping the fixed access/desktop stage and moving directly into the mobile/wireless broadband stage. Active mobile

4 Readers interested solely in the difference between absolute and normalized numbers, rather than the method of measuring these numbers, can skip ahead to the section "Trends in the Vectors of Attack: Vulnerabilities and Malicious Sites."

5 Studying cyberspace from an empirical perspective involves a bit of irony. While we live in the age of big data, where nearly everything a person does online is tracked and recorded, most of this information is proprietary and fragmented among numerous private actors. The result is that it is not easy to get a clear picture of either the size of the Internet or the occurrence of cybercrime. Data, therefore, have to be drawn from multiple sources and often estimates have to be used in place of actual figures. As a disclaimer: all the data used in this paper presents at best a partial view of the actual ins and outs of cyberspace. Despite the fact that many of the sources consulted lay out their data collection procedures, it is not clear how random of a sample of Internet activity the data actually depicts, and so extrapolating from these findings to the entirety of cyberspace can only be done with great care.

Table 1: The Size of Cyberspace

	Minimum	Maximum	Mean	Standard Deviation
Internet Users	1,562,067,594	2,925,249,355	2,252,889,661	500,996,210
Email Users	1,300,000,000	2,504,000,000	1,951,333,333	514,583,586
Active Mobile Broadband Accounts	422,000,000	2,693,000,000	1,318,000,000	808,928,097
Number of Smartphones	139,290,000	1,244,890,000	567,862,857	419,380,858
Number of Domains	177,000,000	288,000,000	230,042,857	41,667,488
Number of Websites	172,338,726	968,882,453	471,754,976	307,845,943
Volume of Data Flows (Gigabytes)	1.2209×10^{11}	7.6685×10^{11}	4.10154×10^{11}	2.46421×10^{11}
Volume of Mobile Data (Gigabytes)	396,816,000	42,336,000,000	13,020,825,714	15,811,807,798
Number of Google Searches	637,200,000,000	2,161,530,000,000	1,538,311,571,429	5.83699×10^{11}
Internet's contribution to GDP (Boston Consulting Group)	1.92×10^{12}	2.45×10^{12}	2.19207×10^{12}	2.18547×10^{11}
Internet's contribution to GDP (McKinsey & Company)	1.42×10^{12}	1.72×10^{12}	1.57879×10^{12}	1.25132×10^{11}

broadband subscriptions are a measure of individuals who access the Internet via a mobile device, such as a smartphone or tablet. They are a smaller, yet rapidly growing, subset of all Internet users. Data on active mobile broadband subscriptions is taken from the International Telecommunication Union's (ITU's) statistics (ITU 2015).

One user can operate multiple devices online (Evans 2011). Each device can potentially be subject to a cybercrime, meaning one person can be targeted multiple times even if one device is only targeted once. Data on the number of smartphones sold to end-users per year is used as a rough proxy for the number of devices online. The number is far, far smaller than the actual number of devices connected to the Web at any one time, but it is likely indicative of the growing trend in connected devices. Data on the number of smartphones sold to end-users is taken from Statista (2015).

Points of Online Interaction

Domains give a good sense of the size of the online ecosystem, as they are a key point of interaction with users. Internet domains include generic top-level domains (such as .com or .net) and country top-level domains (such as .ca and .uk). All domains are registered with the Doman Name System (DNS), which ensures that each domain is globally unique and that when you type in a web address you are taken to the correct website. Data on the number of domains from 2008 to 2014 is taken from Verisign's *Domain Name Industry Briefs* (2008; 2009; 2010; 2011; 2012; 2013; 2014).

The number of websites online is again a good measure of the number of points of interaction online and so a good measure of the size of the Internet ecosystem. There is significant overlap between websites and domains, although the number of websites is larger because one website can have multiple subsidiary pages and because not all websites are actually a part of the DNS. In 2014, the number of websites was 680,882,453 higher than the number of domains. Data on websites is taken from Internet Live Stats (2015b) for the period 2008 to 2014.

Volume of Online Activity

The Internet is essentially a hyper efficient way to send and receive data. Statistics on the volume of data that traverses the Internet, therefore, is a useful measure of how busy the Internet ecosystem is year over year. The Internet is composed of a number of privately run networks that interconnect to provide the system with a global reach (Woodcock and Adhikari 2011). Each network maintains its own records, and piecing together exactly how much data flows globally is extremely difficult. As such, any figure for the size of global data flows is only an estimate. For this paper, data on the volume of Internet traffic from 2008 to 2013 was gathered from the "2009 Cisco Visual Networking Index: Forecast and Methodology, 2008–2013" and data on 2014 was taken from the 2010 iteration of this white paper (Cisco Systems 2009; 2010). The data taken from these reports are Cisco System's estimates of global Internet traffic flows. Despite the best efforts of Cisco System engineers, the data probably under-represent the true size of data flows across the Internet. They also fail to distinguish between the types of data flows (that is, streaming video versus emails and website visits), which could affect the appropriateness of normalizing cybercrime numbers around this metric.

Mobile traffic is a smaller, but rapidly growing, subset of all Internet traffic. Mobile traffic gives a rather obvious impression of how much people are using cyberspace via a mobile device. Mobile operating systems and security systems are distinct from traditional desktop-style systems, with their own weaknesses and vulnerabilities. The volume of mobile traffic shows how much mobile devices

are used to access the Internet and, correspondingly, how likely they are to be the subject of a cybercrime. Data of mobile traffic is also taken from Cisco's two forecasting reports.

The Internet is also, as it is colloquially known, an "information superhighway." Another measure of the activity that occurs on the Internet, therefore, is the number of search engine queries per year. Data on the annual number of Google searches was used as a measure for Internet search queries(Statistics Brain 2015). Globally, Google Chrome is also the largest web browser in every region of the world (StatsCounter 2015). These trends suggest that Google searches are a good proxy for the occurrence of Internet-based searches more generally.

The Internet is becoming increasingly integrated into every aspect of society. One of the most meaningful (or at least most measureable) effects of this growing integration and importance is the Internet's share of global GDP. Currently, no comprehensive time series data exists for this measure. To operationalize the Internet's contribution to global GDP, two separate estimates on the Internet's contribution to various nations' GDP are used here. First is a McKinsey & Company estimate on the contribution of the Internet to the economy of 13 large nations in 2009.[6] Together, these 13 nations make up some 70 percent of the world's GDP. Although the Internet's contribution to global GDP is likely larger than outlined in the McKinsey & Company study, the findings are fairly indicative of the Net's general effect on global GDP. The second measure for the size of the global Internet economy is from a Boston Consulting Group study that looks at the Internet's contribution to GDP in Group of Twenty (G20) nations in the year 2010 (Dean et al. 2012). Together, the G20 makes up around 70 percent of the world's population and close to 90 percent of Global GDP (Griffith-Jones, Helleiner and Woods 2010, 25). Again, the Boston Consulting Group's study provides a partial, but still strongly indicative, picture of the Internet's contribution to global GDP. On average, and this is important to note for the later analysis, the Boston Consulting Group's 2010 estimates of the Internet's contribution to the global economy are, as one would expect, larger than the McKinsey & Company's estimates for the size of the Internet's contribution in 2009. This is in line with the rather intuitive idea that the Internet's contribution to the global economy is becoming proportionately more important over time. The Boston Consulting Group's figures are also more representative of the global contribution of the Internet because they include more countries. As such, even though the McKinsey & Company and the Boston Consulting Group estimates point to similar patterns vis-à-vis the absolute numbers, this paper relies on the more inclusive estimates of the latter in the analysis below.

One additional assumption involving the GDP numbers needs to be laid bare. Both studies provide only a static snapshot of the Internet's contribution to global GDP, one in 2009 and one in 2010. In using these data in the comparisons below, it is assumed that the Internet's proportional contribution to each country's GDP remains constant, so if, as in the case of Sweden in the McKinsey & Company study, the Internet contributed 6.3 percent to the country's GDP in 2009, it is assumed that it also contributed 6.3 percent in 2008 and will only contribute that amount moving forward from 2009 into 2013. Since the Internet and Internet-enabled platforms are becoming increasingly common in business, industry and commerce, this assumption likely works against the real world trend of the Internet expanding in its importance to the economy year over year. The assumption is necessary, however, to get enough data in normalize cybercrime trends against an indicator of the economic size and importance of the Internet. This assumption will effectively under-represent the growing size of the Internet economy and thus shrink the denominator in the normalization of cybercrime statistics below. The assumption (although needed) will paint a picture of the security of cyberspace that is likely worse than what actually exists.

THE SECURITY OF CYBERSPACE: VECTORS, OCCURRENCE AND DAMAGE

The security of cyberspace can be conceptualized best from a user's perspective, broadly defined. A secure cyberspace is one in which a user can make use of the Internet without an unreasonable fear of suffering a high cost, with cost being defined in some combination of reputational, monetary and rights violations terms. An insecure cyberspace environment is the opposite, or basically one in which using the Internet is likely to impose a large cost upon the user. This section outlines how to operationalize the level of security in cyberspace by looking at the available vectors for attack, the occurrence of online cyber attacks and the costs of successful attacks. Together, these three categories give a sense of how insecure cyberspace is for an individual user.

Many aspects of the security of cyberspace are worsening over time, but many others are actually remaining fairly static year over year. In the odd case, a given indicator is actually improving. These measures of the insecurity of cyberspace are akin to the crime rate in a city or country. If they are increasingly slower than the population, staying the same size as the population grows, or improving as the population increases, the common result is an improved crime rate.

6 The countries included in the McKinsey study are Sweden, the United Kingdom, South Korea, Japan, the United States, Germany, India, France, Canada, China, Italy, Brazil and the Russian Federation (Pélissié du Rausas et al. 2011).

Table 2: Summary Statistics for the Security of Cyberspace

	Minimum	Maximum	Mean	Standard Deviation
New Vulnerabilities	4,814	6,787	5,749	781.880
Malicious Web Domains	29,927	74,000	53,317	13,769.99
Zero-day Vulnerabilities	8	24	14.85714	6.336
New Browser Vulnerabilities	232	891	513	240.570
Mobile Vulnerabilities	115	416	217.35	120.85
Botnets	1,900,000	9,437,536	4,485,843	2,724,254
Web-based Attacks	23,680,646	1,432,660,467	907,597,833	702,817,362
Average per Capita Cost	188	214	202.5	8.893818078
Organizational Cost	5,403,644	7,240,000	6,233,941	753,057
Detection and Escalation Costs	264,280	455,304	372,272	83,331
Response Costs	1,294,702	1,738,761	1,511,804	152,502.2526
Lost Business Costs	3,010,000	4,592,214	3,827,732	782,084
Victim Notification Costs	497,758	565,020	523,965	30,342

This conceptualization of the security of cyberspace can be expressed as a function of three factors:

- the vectors available for cyber attack;
- the occurrence of cyber attacks; and
- the damage caused by successful cyber attacks.

Together, these three factors determine how secure cyberspace is for an individual user. For instance, when the vectors of attack are few, cyber attacks are harder to effectively launch, making the cyberspace environment more secure. When the number of attacks is low, the probability that a user will be subject to a cyber attack is less, again making cyberspace more secure. Likewise, when the damage caused by a successful attack is low, the cost of a successful cybercrime for an individual is less severe, meaning the environment is less threatening overall. In every case, as the vectors, occurrence or damage of cyber attacks goes up, the overall security of cyberspace from a user's perceptive goes down.

This paper operationalizes the concept of the vectors of cyber attack via the following measures:

- new vulnerabilities;
- malicious web domains;
- zero-day vulnerabilities;
- new browser vulnerabilities; and
- mobile vulnerabilities.

The concept of the number of attacks are operationalized via:

- botnets; and
- recorded web-based attacks.

And the concept of the damage of attacks are operationalized via:

- average cost per data breach;
- overall organizational cost from data breaches;
- the cost of detecting a data breach and escalating;
- post-breach reaction costs;
- lost business costs; and
- victim notification costs.

Table 2 presents some basic summary statistics on the various indicators of the insecurity of cyberspace.

Vectors of Attack

New vulnerabilities are exploitable points in the software code underwriting a program that can provide a cybercriminal with unwanted access to a device.[7] New vulnerabilities are distinct from zero-day vulnerabilities in that they are publicly known. Companies provide routine updates to their programs (Microsoft updates roughly every Wednesday, for example). These updates often include patches for newly discovered vulnerabilities. Failure to update a program can lead to serious problems, as cybercriminals can exploit peoples' sluggish behaviour to infect a system through these publicly known, but inadequately patched, weak points. Data on new vulnerabilities from 2008 to 2014 are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (Norton Symantec 2009; 2010; 2011; 2012; 2013; 2014; 2015).

7 In the case of the various vulnerabilities discussed in this paper, the numbers are a count of the new vulnerabilities for that year and not a count of all the vulnerabilities that have ever been discovered.

Malicious web domains are domains that have known bits of malicious code embedded within them. This code is designed to infect a visiting user's computer with a virus. Malicious web domains are a passive vector of attack for cybercriminals because they require that the user go to an infected domain. Nevertheless, this can still be a potent avenue of attack. Data on malicious web domains are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

New zero-day vulnerabilities are vulnerabilities in software code that are as of yet unknown. The "zero day" part of the name refers to the fact that there have been zero days available to provide a patch that fixes the vulnerability. Zero-day vulnerabilities are fairly rare and quite valuable. Cybercriminals that gain access to a zero-day vulnerability can attack computers easily, as there is no defence against this exploitation; therefore, they are a highly potent vector of attack. Data on zero-day vulnerabilities are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

New browser vulnerabilities are weak points in the code of web browsers, such as Google, Safari and Internet Explorer. As most of the top level of the Internet is digested via a web browser, they are useful avenues for attack by cybercriminals. The data on web browser vulnerabilities are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).[8]

New mobile vulnerabilities refer to vulnerabilities that are specific to mobile devices, such as Android devices or iPhones, rather than laptops and desktop computers. The data on mobile vulnerabilities are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

Occurrence of Cyber Attacks

Botnets are computers that have been infected by a virus that allows them to be hijacked and used remotely by a third party for some illicit purpose. Botnets are often employed in distributed denial of service (DDoS) attacks, which require that a large number of requests be made of a website in a short period of time. Botnets are also often used to send spam emails. To become a part of a botnet, an online device needs to have been the subject of a cyber attack. A measure of botnet computers is one way to get at the number of victims of a crime, although certainly not the only one. The number of botnet computers, therefore, gives a sense of the occurrence of successful cyber attacks. Data on botnets are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

Recorded web-based attacks are cyber attacks that were launched against one part of the network from an online source and are a good measure of the occurrence of cyber attacks. These attacks exclude cyber attacks that result from, say, the use of an infected USB key. Web-based attacks provide a picture of the overall occurrence of cyber attacks, although, due to reporting problems and the fact that cybercriminals often try to have their attacks go unnoticed, the actual number of attacks is probably higher than the recorded figure. Data on web-based attacks are drawn from the IT security firm Kaspersky Lab's "Security Bulletin" reports (Kaspersky Lab 2008; 2009; 2010; 2011; 2012; 2013; 2014).

The Damage of Cybercrime

The concept of the damage done by cybercrime is operationalized in five ways. This paper focuses exclusively on the cost of data breaches for companies, although this is certainly not the be-all and end-all of the costs cybercrime imposes on to users of the Internet. All the data on breaches is taken from the Ponemon Institute's annual *Cost of Data Breach Study*, which records the overall cost of data breaches (Ponemon Institute 2011; 2013; 2014). Unfortunately, the Ponemon Institute only started collecting a global sample in 2013 and previously only collected the costs associated with US data breaches. The United States is still in the later global assessments, so for the purpose of over-time comparability, only the US numbers are included in the analysis below. Due to the overall lack of statistics on data breach costs, this paper makes the assumption that the US cost of cybercrime data is indicative of the world's costs. In reality, the average costs for the world are almost certainly far lower than the US costs. For example, in 2013, the organizational cost of data breaches in the United States was US$5,850,000. Globally, the average based on the weighted numbers from the later Ponemon Institute studies, including the United States, is US$2,282,095, or a difference of over twice as much. Using the US numbers, in other words, will overstate the costs of cybercrime and actually work against the argument herein that the security of cyberspace is better than the impression given by the absolute numbers.

Before turning to a discussion of the various measures used to operationalize the cost of cybercrime, it is important to note two additional limitations to the statistics collected on data breaches. The companies studied vary from year to year, as does the number of companies that are observed. Clearly, from a methodological point of view, this is not ideal, as the shifting foundational sands of the studies means that the inter-year samples are not

8 There is a major discrepancy in the Norton 2013 report compared to the Norton 2014 report. The 2013 report indicates that new browser vulnerabilities rose from 591 in 2011 to 891 in 2012 before falling to 351 in 2013. The 2014 report indicates that new browser vulnerabilities rose from 351 in 2011 to 891 in 2012 before declining to 591 in 2013. The paper retains the earlier, 2013, data because it actually works against the hypothesis that the security of cyberspace is better than the absolute numbers by moving a higher number earlier in time. In the tests below, using the 2014 data only changes the magnitude, and not the direction, of the relationship.

strictly comparable. Another limitation is that the studies exclude "mega breaches," or those involving more than 100,000 breached records in a single attack. This restriction essentially excludes high-damage but low-probability events in favour of the more representative high-probability but comparatively low-damage events that occur most of the time. Despite all these limitations, the Ponemon Institute's studies of the cost of data breaches are the best publicly available data on the overtime costs of data breaches.

The first operational measure of the cost of cybercrime is the average cost for a company per breached record. This measure shows the organization's cost divided by the number of compromised files. This measure is one way to show how much an organization has to pay as a consequence of cybercrime.

Another way to portray this cost — and the second measure of the costs of cybercrime — is the overall average organizational cost of data breaches in a given year. This figure is basically the total price tag of dealing with data breaches. It is a good measure of the cost of cybercrime because it quantifies the absolute cost that a company needs to pay as a result of online criminal behaviour.

A third measure of the costs of cybercrime involves a company's detection and escalation costs. Data breaches are bad; undetected data breaches are worse. Companies invest considerable resources into IT security so that they can detect data breaches, and, if warranted, act to repel them, although these sums are not necessarily sufficient. This is a good measure of the cost of cybercrime because it involves the investment that companies need to undertake since they operate in an environment with less than perfect security.

A fourth measure is the cost that an organization needs to pay after a data breach in order to fix any damage done. Cybercrime can often result in damage to software and computer hardware. This is a good measure of the cost of cybercrime, because, like a broken window after a burglar breaks into a person's home, the damage done by cybercrime is not just a result of what is stolen.

A fifth measure of the costs of cybercrime is the cost of lost business. Companies, in particular those that provide an online service, rely on the public's perception that their services are trustworthy. If the public thinks that using a company's services will lead to a loss of personal or financial information, individuals are likely to choose other service providers or cease that activity entirely. The cost of lost business as a result of the occurrence of data breaches is a good measure of the sort of second-order effect of cybercrime on a company's balance sheet.

A final measure of the costs of cybercrime is the cost of notifying victims that their records, be they personal, financial or otherwise, have been compromised in a data breach. Even though companies might have an incentive to cover up a data breach for fear of losing business, many are legally obliged to inform those individuals that have had their information compromised.

TRENDS IN THE VECTORS OF ATTACK: VULNERABILITIES AND MALICIOUS SITES

This section compares the absolute numbers for the various vectors of attack against the normalized trend. In every case, the normalized trends presents a picture of the security of cyberspace that is better than the one presented by the un-normalized absolute figures.

This section looks at vectors of cyber attack, which are basically the ways in which cyber attacks can occur to an Internet user. The relative number of ways in which an Internet user can be attacked are declining, given the growing size of the Internet. One way to think of this is to imagine a city with a number of high-crime neighbourhoods. If the city is made up of 10 neighbourhoods and five of them are dangerous, then the crime rate is 50 percent. If the city grows (as cyberspace has grown) faster than the number of bad neighbourhoods, then the crime rate declines and people are relatively safer. Imagine the hypothetical city grows in size to 15 neighbourhoods, but the number of high-crime areas stays at five. The new crime rate is only 33 percent. The city is safer as a result and a person's chance of being subject to a crime declines. Cybercrime vectors are like the high-crime neighbourhoods.

The analysis below undertakes a number of different normalizations for each measure of the security of cyberspace. A justification for each normalization is provided in each section. Multiple normalizations are used, rather than just a single one for each measure of cybercrime, because there is not an agreed-upon denominator that makes the most sense across the different measures. So, in the interest of painting the broadest possible picture and of forestalling the notion that this paper uses only the normalizations that support its argument, several normalizations per cybercrime measure are included.

Figure 1 normalizes new vulnerabilities as a vector of attack around the number of Internet users, the number of email users and the number of websites. Since vulnerabilities are weaknesses in computer code, the ideal denominator for new vulnerabilities would be the number of software programs that are in use around the world. Unfortunately, the number of programs is not even partially known. In the absence of this data, Internet users, email users and websites will have to do. The number of Internet users gives an (admittedly partial) impression of the number of devices that are operating online and so indicates the

chance that a device will be using software that is afflicted by a new vulnerability. The number of email users is another measure of active devices online, pointing to the odds that a device will be running a flawed program. Finally, websites are hosted using various software programs, all of which can have unexpected vulnerabilities. The number of websites, therefore, provides a measure of the points of interaction online that are operating software that could be prone to cyber attack due to a new vulnerability.

In Figure 1, the trend in the absolute figures suggests that the number of new vulnerabilities is actually worsening between 2008 and 2014, rising from 5,562 new vulnerabilities in 2008 to 6,549 new vulnerabilities in 2014; an increase of 17.75 percentage points over the five years. In contrast, each of the normalized trends suggests that this vector of attack is actually improving over time. For instance, new vulnerabilities normalized around the number of Internet users, a proxy for online devices in this case, fell from 3.56 new vulnerabilities per 1,000,000 Internet users in 2008 to 2.24 vulnerabilities per 1,000,000 Internet users in 2014. This drop amounts to a percentage change of 37.13 percent. In other words, the normalized numbers suggest that the security of cyberspace is greater than what is suggested by the absolute numbers. Indeed, the absolute numbers indicate that the situation is worsening, while the normalized figures actually indicate that the situation is improving.

Figure 1: New Vulnerabilities

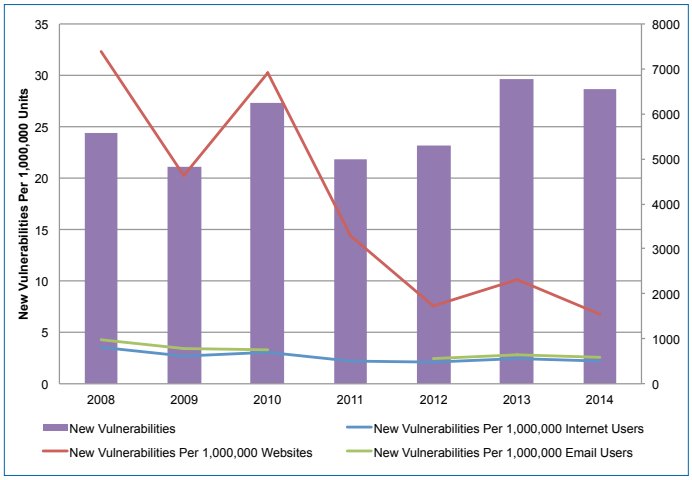

Figure 2 compares the normalized trend among malicious domains as a vector of attack against the absolute number of malicious domains. The number of malicious web domains is normalized around the number of Internet users, the number of web domains and the number of websites. Clearly, the most natural data manipulation is to normalize malicious domains around the total number of domains (which is done in both Figure 2 and then in more detail in Figure 3). Normalizing malicious domains around the number of Internet users makes sense because the latter measures the number of people that can be affected by a malicious domain, which shows the trend in potential infection rates. As mentioned above, the number of web domains is a smaller subset of the total number of websites, which can have subsidiary pages and the like. Normalizing the number of malicious domains around the number of websites provides another glimpse of how problematic a given number of malicious web domains are likely to be because it shows how many websites might be affected and so how many webpages might be a threat to the security of cyberspace.

As shown in Figure 2, the number of absolute new malicious web domains has remained fairly constant over time, with an initial increase from 2010 to 2012 being followed by a decline from 2012 to 2014. In contrast to these fairly stable numbers, the normalized trends in malicious web domains per 1,000,000 Internet users and per 1,000,000 websites both strongly point toward an improving security situation in cyberspace. However, probably the most appropriate normalization in this case is the number of malicious web domains per 1,000,000 Internet domains, since the basic unit of measure (domains) is the same. Here, the absolute number of malicious domains and the normalized trend track together fairly consistently, but the actual trend underlying the two sets of data shows a clear difference in degree.

Figure 2: New Malicious Web Domains

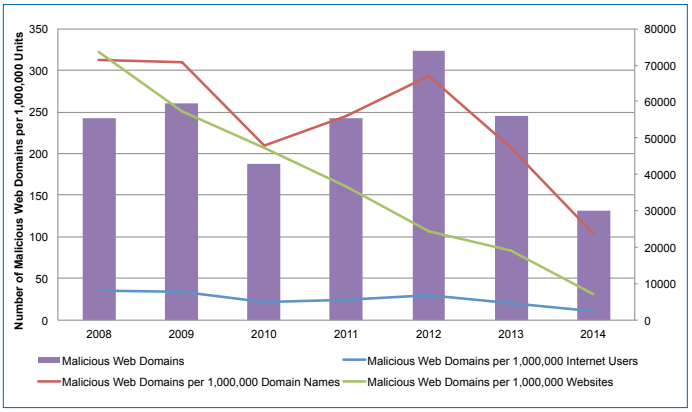

Figure 3 looks at just the comparison of the absolute number of malicious domain names and the trend in malicious domains normalized around the total number of domains. The appearance that these two indicators track together over time suggests that there is a fairly static proportion of all web domains that are malicious. However, this initial impression is misleading in the sense that the two sets of numbers are changing at very different speeds. The two trend lines in Figure 3 show that between 2008 and 2014 both the absolute and the normalized trends have been improving. Comparing the rate at which the situation is improving tells a different story. The absolute number of new malicious domains has fallen from 55,389 malicious domains in 2008 to 29,927

malicious domains in 2014, a decline of 45.96 percent. In contrast, the normalized numbers fell from 312.93 malicious domains per 1,000,000 domains in 2008 to only 103.91 malicious domains per 1,000,000 domains in 2014, which amounts to a decline of 66.79 percentage points. As with the new vulnerabilities, the data from Figures 2 and 3 support the idea that the absolute numbers overrepresent the insecurity of cyberspace compared to the normalized trends by showing the picture improving more slowly than is actually the case.

Figure 3: Normalized versus Absolute Domains

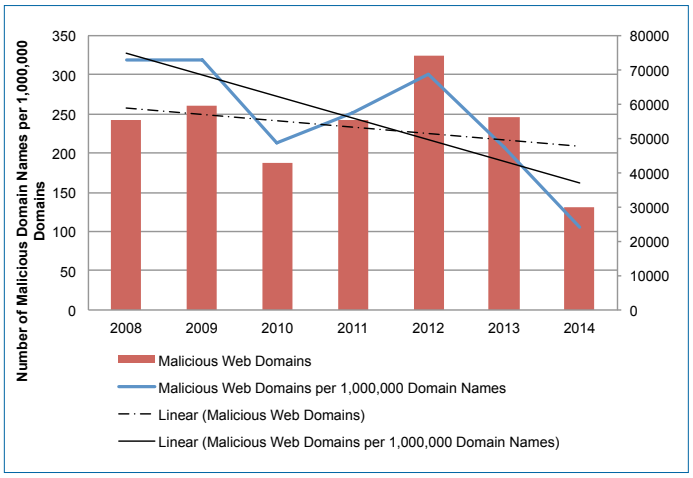

Figure 4 presents the data on the number of zero-day vulnerabilities normalized around the number of Internet users, web domains and the number of websites, and contrasts these numbers with the absolute trend. As with new vulnerabilities, the best measure to normalize zero-day vulnerabilities around would be the number of software programs used in the world, the data for which does not exist. Nevertheless, since zero-day vulnerabilities are weaknesses in computer code, the normalization that makes the most sense is the number of zero-days per 1,000,000 websites, since websites rely on a growing number of software platforms (think of the Heartbleed zero-day exploit in Secure Sockets Layer [SSL] in 2014). In the interest of presenting the broadest possible story, the number of zero-day vulnerabilities normalized around the number of Internet users and email users are also included (both proxies for the number of potentially vulnerable devices operating various pieces of software).

The dotted trend line in Figure 4 shows that over time the absolute number of zero-day vulnerabilities is getting larger, suggesting a worsening cyber security environment. This finding is mirrored by the trend in zero-day vulnerabilities per 1,000,000 email users and per 1,000,000 Internet users. However, the trend in zero-day vulnerabilities per 1,000,000 websites is actually declining over time, despite a jump upward in 2013. To the extent that normalizing the number of zero-day vulnerabilities around the number of online websites is the most accurate measure of this vector of cyberattack, the fact that the trend is negative suggests that, as is the case with the other measures, the security of cyberspace is improving over time even as the absolute number of zero-day exploits increases.

Figure 4: New Zero-day Vulnerabilities

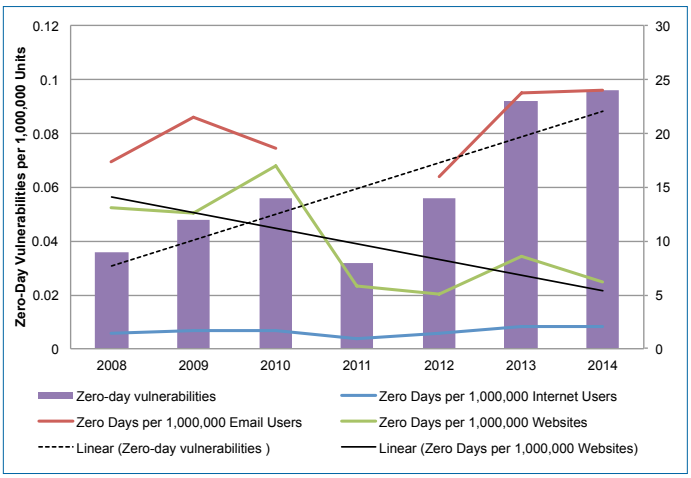

Figure 5: New Browser Vulnerabilities

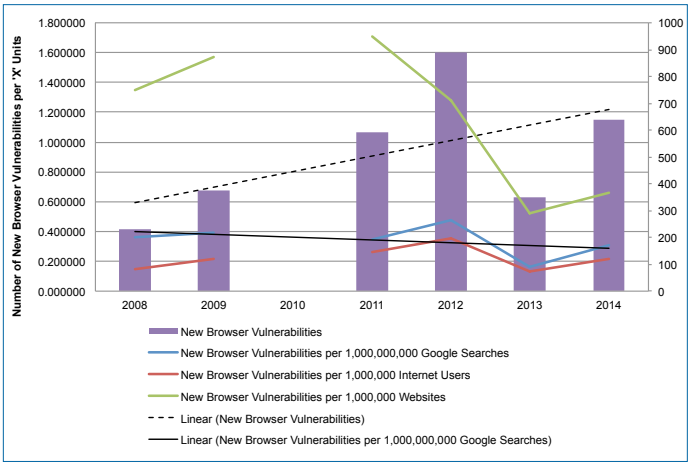

Figure 5 summarizes the data on browser vulnerabilities as a vector of cyber attack, depicting both the absolute numbers and the number of new browser vulnerabilities normalized around the number of Internet users, the number of websites and the number of Google searches. The number of new browser vulnerabilities are normalized around the number of Internet users because this manipulation of the data shows the rate at which people will come into contact with vulnerable browsers (not accounting for the fact that different browsers are used more frequently than others). The number of new browser vulnerabilities are normalized around the number of websites because these are the points of online interaction that people are trying to reach via a web browser. The more websites that exist, the more people will be pulled to use a web browser and so the larger the potential that a browser will affect an online device. Finally, in what is

probably the most accurate normalization, the number of browser vulnerabilities is divided by the number of Google searches. Google searches capture the frequency with which a globally dominant web browser is actually being used and thus how probable it is that an Internet user will come into contact with a vulnerable browser.

As shown by the dotted trend line in Figure 5, the absolute number of new browser vulnerabilities is generally increasing over time, with 639 browser vulnerabilities in 2014 compared to 232 in 2008 (an increase of 175 percentage points). New browser vulnerabilities normalized around the number of Internet users is also slightly escalatory over the full seven-year period. In contrast, new browser vulnerabilities as a proportion of all websites shows a generally de-escalatory trend and an improving cyber security situation. Most telling, given its likely accuracy as a measure of effect of new browser vulnerabilities, the number of vulnerabilities normalized around Google searches is negative, as shown by the solid black trend line. In numerical terms, the number of new browser vulnerabilities per 1,000,000,000 Google searches drops from 0.364 new vulnerabilities per 1,000,000,000 Google searches in 2008 to 0.305 new vulnerabilities per 1,000,000,000 Google searches in 2014, a decline of 16.23 percentage points. Overall, the numbers on new browser vulnerabilities as a vector for cyber attack again support the idea that the absolute numbers paint a worse picture of the security of cyberspace than the normalized numbers. In this case, the absolute numbers indicate that the situation is worsening, while the normalized numbers say that things are actually improving.

Finally, Figure 6 shows the number of new mobile vulnerabilities and the number of new mobile vulnerabilities normalized around the number of active broadband mobile subscribers, the number of smartphones sold to end-users, and the volume of mobile data usage in gigabytes. These three normalizations make eminent sense because mobile vulnerabilities (glitches and weaknesses in the operating system or associated software of mobile devices) can only affect mobile users. Each normalization helps clarify the real risk that a user faces when using a mobile device to access the Internet. Normalizing new vulnerabilities around active mobile broadband subscriptions shows how likely a user is to be affected by a new vulnerability. Normalizing the number of new vulnerabilities around the number of smartphones sold to end-users shows the likelihood that a particular device will be afflicted by a cybercrime. Finally, normalizing the number of new mobile vulnerabilities around the volume of mobile traffic shows how problematic weakness are in light of how much people use mobile platforms to access the Internet.[9]

As shown in Figure 6, mobile vulnerabilities have expanded rapidly since 2009, with the number of new mobile vulnerabilities increasing from 115 in that year to 415 at the peak in 2012, before declining to 127 in 2013 and jumping up again to 168 in 2014. This growth in mobile vulnerabilities tracks the growth in the use of mobile devices, both in the developed world and among new entrants to the Internet. From 2009 to the peak (in terms of new mobile vulnerabilities) in 2012, the absolute numbers indicate that the number of new vulnerabilities rose by 261 percentage points. Across the whole sample, the absolute numbers on new mobile vulnerabilities indicate that the security of cyberspace is growing worse over time, even with the significant drop in new vulnerabilities in 2013, as shown by the long-dashed trend line. In contrast, the three normalized measures each show that the security of cyberspace is actually improving. The reduction in new vulnerabilities relative to the various measures is also substantively large. For example, the number of new vulnerabilities per 1,000,000 gigabytes of mobile data fell from 0.29 vulnerabilities per 1,000,000 gigabytes in 2009 to 0.0064 vulnerabilities per 1,000,000 gigabytes in 2014, a reduction of roughly 97.7 percentage points. Active mobile broadband subscriptions, for their part, fell from 0.273 new vulnerabilities per 1,000,000 subscriptions in 2009 to 0.086 vulnerabilities per 1,000,000 subscriptions in 2014, a reduction of 68.43 percentage points. Finally, the number of new vulnerabilities per 1,000,000 smartphones sold fell from 0.826 in 2009 to 0.173 in 2013, a reduction of 79.02 percentage points. Clearly, the normalized numbers paint a radically different picture of the security of cyberspace than the absolute numbers, the latter showing the situation getting worse and the normalized numbers showing the situation rapidly improving. In short, mobile vulnerabilities continue to grow, but they are growing more slowly than the actual use of mobile devices. Essentially, the absolute numbers say that the situation is worsening, when, as shown by the normalized numbers, the security of cyberspace is actually improving.

9 Clearly, the best measure in this case would be if both vulnerabilities and broadband subscriptions specified the type of operating system or software that was problematic and used on the device. Since this data does not exist, the data included in the text is the next best option.

Figure 6: New Mobile Vulnerabilities

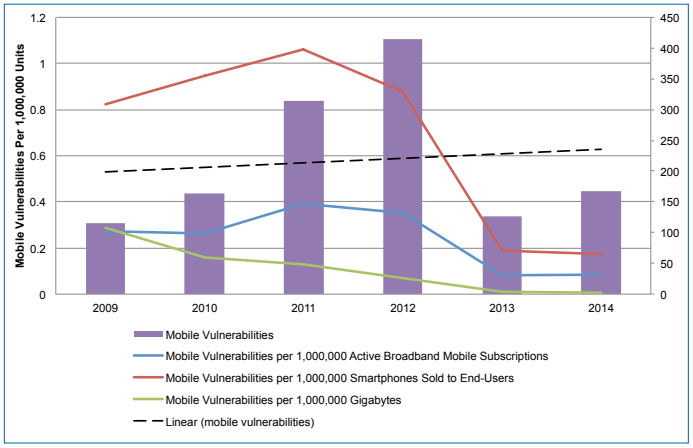

When it comes to the potential vectors of cyber attack, the security of cyberspace is far better than what is shown by just looking at the absolute numbers. In four of the five vectors of attack (new vulnerabilities; zero-day exploits; browser vulnerabilities; and mobile vulnerabilities), the absolute numbers say that the situation is getting worse over time, while the normalized numbers show the opposite: cyberspace is becoming more secure. In the remaining case (malicious domains), both the absolute and the normalized numbers indicate an improving situation, but the former shows cyberspace getting better at a slower rate than the latter. In short, when it comes to vectors of attack, cyberspace is a lot safer than one might think.

OCCURRENCE OF CYBER ATTACKS: WEB-BASED ATTACKS

This section looks at the occurrence of cyberattacks in absolute terms compared to the normalized trend in the number of botnet computers and cyber attacks between 2008 and 2014, given the growing size of cyberspace. On botnets, or computers that have been successfully targeted by a cyberattack, both the absolute and the relative numbers show that things are improving over time. The normalized numbers, however, point to a situation that is getting better faster, when compared to the absolute numbers. Both the absolute and the normalized numbers for the occurrence of cyber attacks indicate that the situation has worsened overall since 2008-2009. At the same time, both sets of numbers show the situation improving since 2013 (in the case of the absolute numbers) and 2012 (in the case of the normalized numbers). Yet, the normalized numbers not only show the situation getting better sooner, but also indicate that things are getting better faster, when the growing size of cyberspace is taken into account. Looking at the actual occurrence of cyber attacks, in other words, the absolute numbers again paint a worse picture of the trends than the relative ones.

The occurrence of cyber attacks is like the occurrence of robbery or violent crime in the real world. Cyber attacks directly target Internet users in some way or another, as crime does in the physical space. To be rather selfish about it, you might not really care how much violent crime there is in a city, only your chances of being the subject of that crime. The basic story in cyberspace is that there has been an increase in violent crime in our hypothetical city of 100,000 people since 2008. But, since the early 2010s, the situation has stabilized and even started to improve overall. More pointedly, a person's chances of being the subject of a cybercrime have declined as the size of cyberspaces has grown and the number of attacks has fallen. Things are getting better, even if the golden age of low crime levels seems to have passed.

Figure 7 plots out the absolute number of botnets compared to the number of botnets normalized around the number of Internet users, active mobile broadband subscriptions and email users. These three measures of the size of the Internet mesh well with the nature of botnets. Botnets are hijacked computers, which today can be desktops, laptops, phones, fridges or any other connected device. Once commandeered, these devices can be used to send spam and launch DDoS attacks. To become part of a botnet, a computer needs to become infected with a malicious program. This means that the computer needs to be operational (Internet users, active mobile broadband subscriptions and email users express the number of operational computers, although the number in each case is smaller than the actual number of online devices) and need to be infected somehow (Evans 2011).[10] As such, the three normalizations that make the most sense are botnets divided by online users.

Figure 7: Botnets

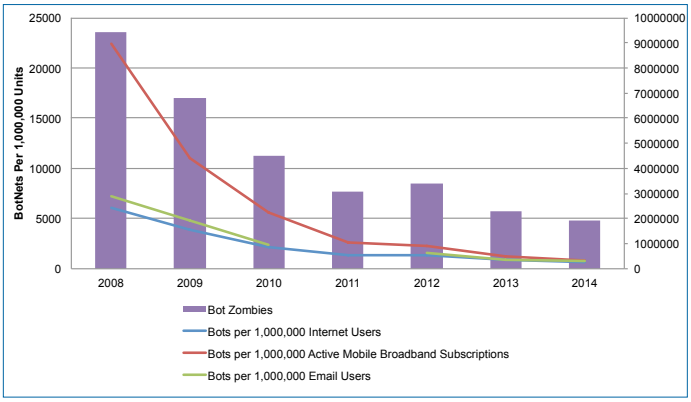

As is clear from Figure 7, while both the normalized and the absolute numbers point to a decline in the number of

10 This conceptualization focuses on the risk of having a computer become a botnet and not the other side of the issue of whether a botnet will be used to launch a DDoS attack on a website. Looking from this angle, the normalization of botnets around the number of Internet, active mobile broadband subscriptions or email users expresses how large the criminal element is as a proportion of all users.

botnet computers between 2008 and 2014, the normalized numbers show a far steeper drop.[11] The absolute number of botnet zombies, which is a count of the number of infected computers worldwide, fell from 9,437,536 in 2008 to only 1,900,000 in 2014, which is a drop of 79.9 percentage points. In contrast, the number of botnets normalized around the number of Internet users fell from 6,041.69 botnets per 1,000,000 Internet users to 650 botnet computers per 1,000,000 users during this same period, amounting to a decrease of 89.24 percent. Similar magnitude declines are found for both active mobile subscriptions (–96.3) and email users (–89.5). This data suggests that the absolute figures overrepresent the insecurity of cyberspace compared to the normalized numbers by exaggerating the problem of botnets as a potential vector of cybercrime.

Figure 8: Web-based Attacks

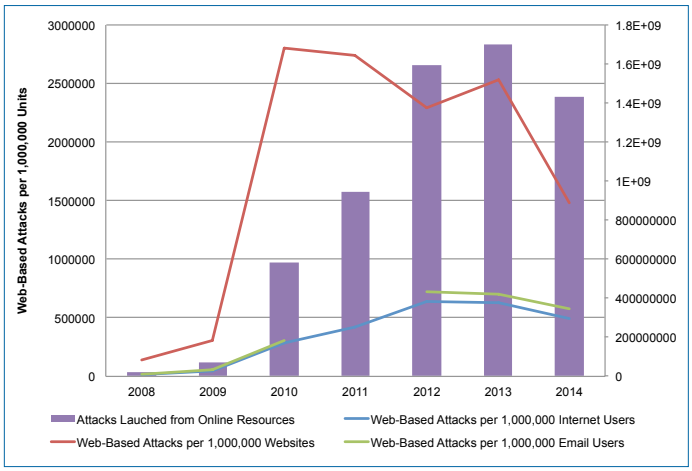

Figure 8 shows the level of absolute web-based attacks compared to the number of such attacks normalized around the number of Internet users, the number of websites and the number of email users. The normalization of the occurrence of attacks around both Internet users and the number of email users captures the idea that cyber attacks target individuals who use the network and that one's chance of being affected by a cybercrime is determined by both the number of attacks and the number of other Internet users. These normalizations, in other words, are similar to normalizing crime statistics around the number of people that live in an affected area. Websites are one clear source of web-based attacks. The normalization of the number of attacks around the number of websites (crudely) shows how frequently attacks occur given the available stock of online points of interaction.

As shown in Figure 8, the absolute numbers point to a strong escalatory trend in cyber attacks, indicating a worse level of security in cyberspace between 2008 and 2014. For example, there were 23,680,646 web-based attacks in 2008 and some 1,432,660,467 attacks in 2014, which is a 5,950 percentage point increase over just seven years! In contrast, the number of web-based attacks per 1,000,000 Internet users has only increased from 15159.8 in 2008 to 489,756.7 in 2014, which is an increase of only (using that term very loosely) 3,130.63 percent. The normalized trends also all suggest that, while the cyberspace security situation is definitely worse than in 2008 and 2009, the trend in normalized cyber attacks has improved since 2010 in the case of attacks per 1,000,000 websites, and since 2012 in the case of attacks per 1,000,000 Internet and 1,000,000 email users. The absolute numbers suggest that, at best, the situation started to improve only in 2014, although it is possible that the low number of web-based attacks in 2014 is a statistical fluke rather than the start of a real trend in the absolute numbers.

Figure 9: Web-based Attacks and Internet Traffic Flows

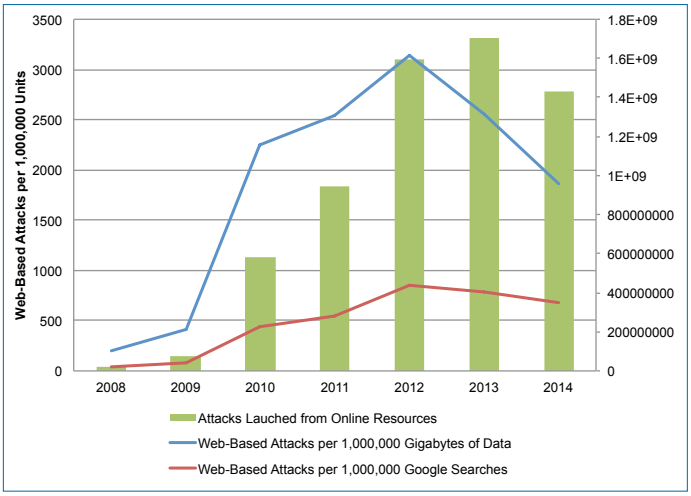

Figure 9 normalizes the number of cyber attacks around the volume of Internet traffic and the number of Google searches for the 2008–2014 period. The intuition behind both normalizations is that, even if there is a constant rate of web-based attacks, the absolute number of attacks should grow as the Internet is used more and more in our daily lives. In such a case, more web-based attacks might not mean an individual user is more likely to be subjected to a cybercrime. What matters is the rate at which web-based attacks occur. Normalizing web-based attacks around the total volume of Internet traffic roughly indicates what proportion of Internet activity is actually malicious and aimed at undermining the security of cyberspace. As a caveat, the rapid growth in video streaming likely biases these numbers, as streaming video takes up a lot of bandwidth and does not usually come with the same

11 The processes for identifying and counting botnets have also improved over time, rendering a more accurate picture of the total number of active botnet computers. While it is impossible to know for sure, it is plausible that earlier counts under-represented the number of botnets, which suggests that the decline has been even steeper. I am grateful to Laura DeNardis for pointing this out to me.

level of security risk as generic web surfing.[12] Normalizing the occurrence of web-based attacks around the number of Google searches is another way to get at the rate at which online activity is likely to be marred by cybercrime. In this case, the measure of online activity is imperfect because Google searches are only a significant subset of all search engine queries and do not encompass all online activity.[13]

As shown in Figure 9, both the absolute numbers and the normalized trends point to an overall escalatory situation in the occurrence of cyber attacks between 2008 and the end of 2014. Yet, there is some hope as web-based attacks fell from 1,700,870,654 attacks in 2013 to 1,432,660,467 attacks in 2014. This amounts to a decline of around 15.77 percent. In contrast, these data show that the normalized trends both start to improve sooner (2012 rather than 2013) and fall more sharply than the absolute numbers. The number of web-based attacks as a share of all Internet traffic, for example, falls from roughly 3,143 attacks per 1,000,000 gigabytes of data in 2012 to roughly 1,868 attacks per 1,000,000 gigabytes of data in 2014, which amounts to a decline of 40.55 percent. The number of web-based attacks normalized around the number of Google searches likewise falls from roughly 852 attacks per 1,000,000 Google searches in 2012 to 684 attacks per 1,000,000 Google searches in 2014, or a decline of 19.7 percentage points. In short, looking at attacks as a proportion of data flow and online activity, the security of cyberspace is again improving both sooner and faster than what is shown by the absolute numbers.

There has indeed been a massive increase in the absolute number of web-based cyberattacks since 2008. Yet, while the glory days of 2008 and 2009 might be gone, since 2010–2012, the rate at which web-based cyber attacks have occurred has declined a lot more than you might otherwise think when factoring in the growing size of the Internet. All five normalized trends bear out this claim.

Overall, the findings in this section show that, when compared to the absolute numbers, the various normalized numbers all point to a situation that both starts improving sooner and that improves more rapidly. The security of cyberspace, in other words, is better than one might think looking at just the absolute numbers.

THE COST OF SUCCESSFUL CYBER ATTACKS

This section compares the absolute numbers to do with the various costs of data breaches with the same numbers normalized around the size of the Internet's contribution to the global economy. Underlying this move is the idea that we need to understand the cost of cybercrime relative to the economic benefits that accrue from the Internet. The real concern would be when the costs of doing business are greater than the benefits produced by using the Internet as a platform for communications and commerce, as firms would then opt out of the system. Normalizing the numbers in this way shifts the question from what a firm pays as a result of data breaches to what sort of economic damage is done in general terms by cybercrime compared to the benefits that are generated by the Internet economy. Again, the absolute numbers consistently suggest a worse cyberspace environment than the normalized numbers.

When it comes to the costs of cybercrime, the value added of the Internet is outpacing the costs that Internet-enabled cybercrime imposes on society. In other words, in net terms, having the Internet is still beneficial, even though cybercrime inflicts economic damage. In the daily world, another example of a sort of dual-use system that both generates economic growth and facilitates crime is the global financial system, which can be used to provide loans and transfer funds, but which can also be used to launder money and avoid taxes. At a social level, what matters are net gains, and, in the case of the Internet and cybercrime — as in the case of the global financial system — things are looking pretty good.

Figure 10: Average Cost per Breach

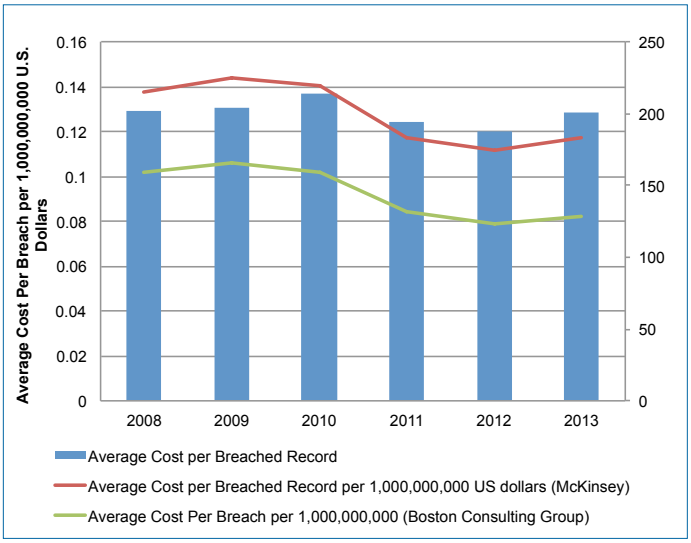

Figure 10 looks at the average cost per breached record in absolute terms compared to these numbers normalized around both the McKinsey & Company and the Boston Consulting Group's estimates for the Internet's contribution

12 I am grateful to the reviewer for pointing out this limitation in the data.

13 A better measure that is not publicly available would be web queries, where people are making requests to view websites. Again, I am grateful to the reviewer for pointing out this potential measure. I only lament that I could not find the data to bring the idea to fruition.

to global GDP. The absolute numbers paint an image of a roughly constant average cost per breached record, with the cost in 2008 being $202.00 and $201.00 in 2013.[14] In contrast, both sets of normalized figures show a reduction over this same time period. The numbers normalized around the McKinsey & Company estimates for how much the Internet contributes to the global economy show a drop in the average cost per breached record from $0.14 cents per $1,000,000,000 dollars of Internet contribution in 2008 to $0.12 cents per 1,000,000,000 in 2013. This decline amounts to a 14.85 percentage change in the normalized cost per breached record. Likewise, the numbers normalized around the Boston Consulting Group estimates show a similar declining trend, with the average cost per breached record per $1,000,000,000 of the Internet's contribution falling from $0.10 cents in 2008 to $0.08 cents in 2010 (a reduction of 19 percentage points). The comparison of the data on the average cost per breached record indicates that the absolute trend depicts a relatively constant level of cost, while the normalized trends show a decreasing cost. Overall, the absolute figures overrepresent the cost of cybercrime in this area compared to the normalized figures.

Figure 11 presents data on the average overall organizational cost that a company is forced to bear as a result of data breaches. A pretty consistent message emerges across all the numbers, with the absolute and normalized trends pointing to a declining cost due to data breaches and thus an overall improvement in the security of cyberspace. However, a comparison of the rate at which the numbers are declining paints a slightly different picture. For the absolute figures, the overall organizational cost fell from a high of $7,240,000 in 2010 to just $5,850,000 in 2013. This drop amounts to a decrease of 19 percentage points. In contrast, looking at the Boston Consulting Group's estimates for the size of the Internet's contribution to GDP, the number falls from a peak value of $3,513.60 for every billion that the Internet contributed to global GDP in 2009 to a low of $2,390.19 per $1,000,000,000 in 2013, amounting to a drop of 32 percentage points.[15] In short, organizational costs due to data breaches are declining across both data forms, but the rate of that decline varies.

Figure 11: Organizational Cost

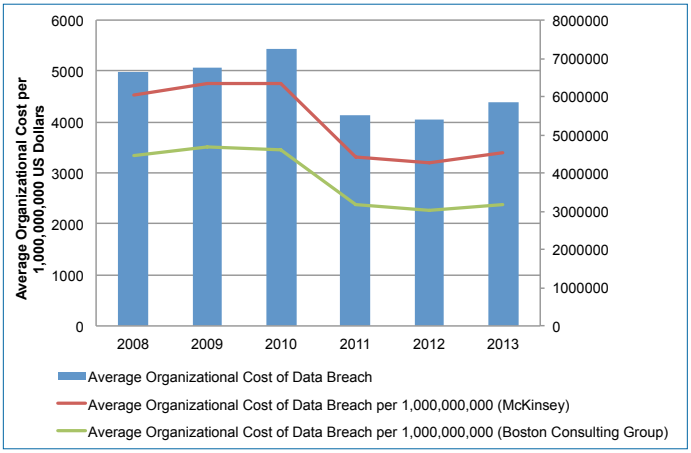

Figure 12 compares the absolute and normalized costs associated with detection and escalation in response to a data breach. In this case, all three sets of numbers point to growing detection and escalation costs since 2008, with a slight reduction in the costs since 2010. Once again, focusing on the magnitude of the changes provides interesting nuance to the picture. From 2008 to 2013, the absolute cost of detection and escalation rose from $271,084 to $417,700 or an increase of 54.1 percentage points. In contrast, the numbers normalized around the Boston Consulting Group estimates of the size of the Internet's contribution to global GDP show that the costs have increased from $136.17 per $1,000,000,000 in 2008 to only $170.66 per $1,000,000,000 in 2013. This change amounts to only a 25 percentage point increase over that time period. In short, the normalized trends show that the growth in the costs of escalation and detection is less pronounced compared to the absolute figures. A similar story of different magnitude changes emerges if we look at the drop from the high point of detection and escalation costs in 2010 compared to the costs in 2013. Here, the absolute values decline from $455,304 in 2010 to $417,700 in 2013, or a decrease of roughly 8.3 percentage points. In contrast, the numbers normalized around the Boston Consulting Group estimates decrease from $216.93 per $1,000,000,000 in 2010 to $170.66 dollars per $1,000,000,000 in 2013, which amounts to a reduction of roughly 21 percentage points.

Overall, the comparison of the absolute and normalized cost of detection and escalation shows that, since 2008, the costs have uniformly increased, but that the absolute numbers have registered a larger percentage increase in that time compared to the normalized numbers. Likewise, since the high point in terms of the costs in 2010, the absolute numbers show a smaller decline in the costs of detection and escalation compared to the normalized trends. Once again, the absolute numbers paint a more dismal picture of the costs of cybercrime than the normalized figures, suggesting that the security of cyberspace is actually greater than is commonly perceived.

14 The addition of a trend line shows a slight decline in the absolute numbers over the full sample.

15 The McKinsey & Company numbers also suggest a larger decline for the normalized trend of around 28.5 percentage points.

Figure 12: Average Detection and Escalation Costs

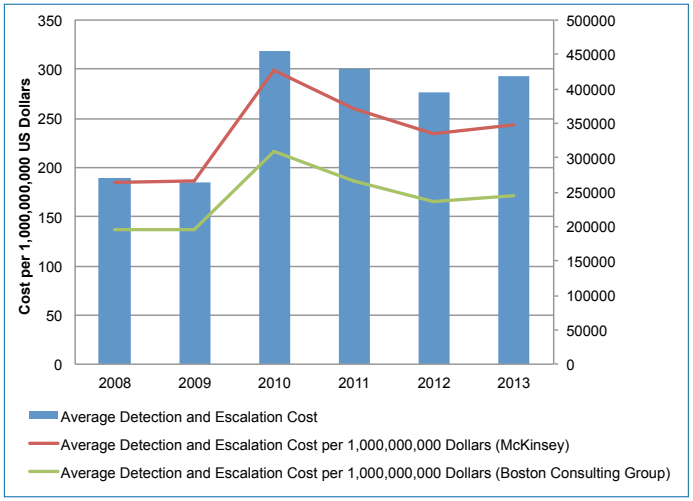

Figure 13 presents data on the absolute and normalized trends in post-breach response costs. At first blush, both the absolute and the normalized numbers paint a roughly consistent picture. A more in-depth comparison reveals two points that suggest the absolute numbers overrepresent the post-breach response costs of cybercrime. First, the absolute numbers indicate an escalatory trend in costs, as shown by the long-dashed trend line. In contrast, the numbers normalized around the Boston Consulting Group estimates for the Internet's contribution to the global economy shows a de-escalatory or declining trend, as shown by the solid black trend line (the McKinsey & Company numbers also point to a declining trend). Secondly, the rate at which the post-breach costs have declined since the high-water mark of 2010 into 2013 shows a greater decline for the normalized numbers compared to the absolute numbers. In particular, the absolute costs fell from $1,738,761 in 2010 to $1,599,996 in 2013 or a decrease of 7.98 percentage points. In comparison, the numbers normalized around the Boston Consulting Group's estimates show a decline from $828.44 per $1,000,000,000 in 2010 to $653.73 per $1,000,000,000 in 2013, which amounts to a decrease of 21.1 percent. With respect to the post-breach response costs, the absolute numbers point to both a worsening situation and a slower rate of potential improvement, while the normalized numbers point toward a generally improving situation and a larger decrease since the highest level of costs in the sample.

Figure 13: Post-breach Response Costs

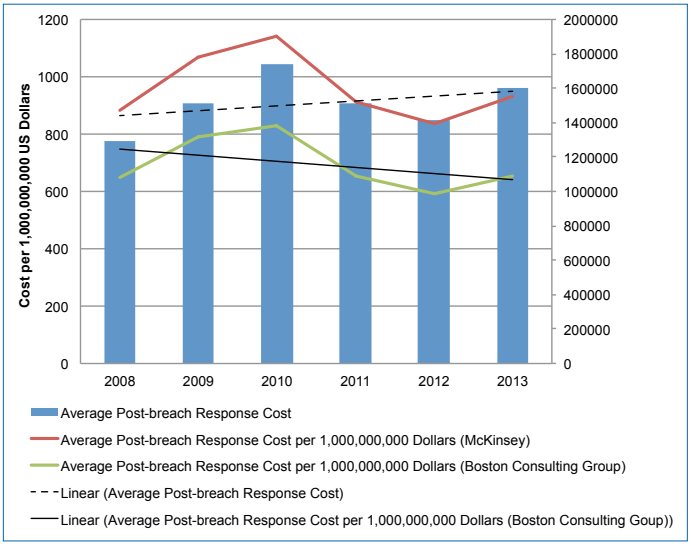

Figure 14: Lost Business Costs

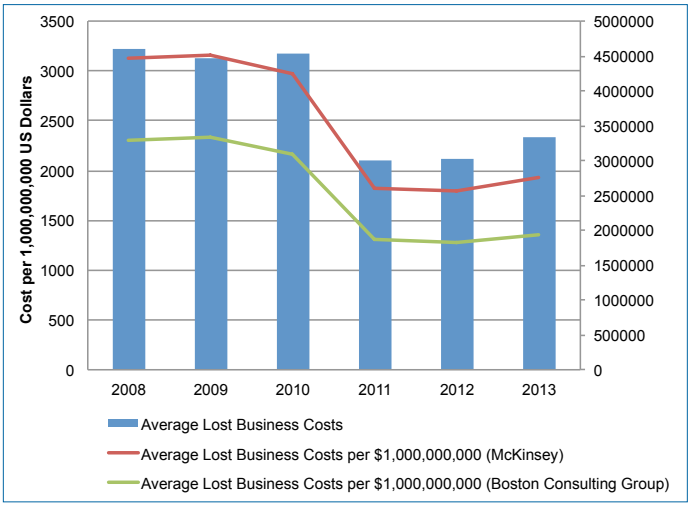

Figure 14 looks at the costs that firms need to endure due to lost business after they have been subject to a data breach. All three sets of numbers show a declining trend in terms of the lost business costs, which could suggest consumers are getting used to data breaches as a part of business in the digital age or that businesses are becoming more adept at managing the public relations side of data breaches. Running a comparison of the rate at which the costs have declined shows again that the absolute numbers depict a comparatively worse environment compared to the normalized trends. For example, in absolute terms, the lost business cost due to cyber attacks faced by firms in 2008 was $4,592,214. By 2013, that number had declined to $3,324,959. The percentage change in the absolute numbers amounts to a decrease of 27.6 percentage points. The numbers normalized around the Boston Consulting Group's estimates for the Internet economy fell from $2,306.74 per billion in 2008 to $1,358.51 per billion in 2013. This change amounts to a decrease of 41.1 percentage

points. Once again, the normalized numbers point to a situation where the lost business costs suffered by firms are improving faster than the costs as they are suggested by the absolute numbers.

Figure 15: Notification Costs

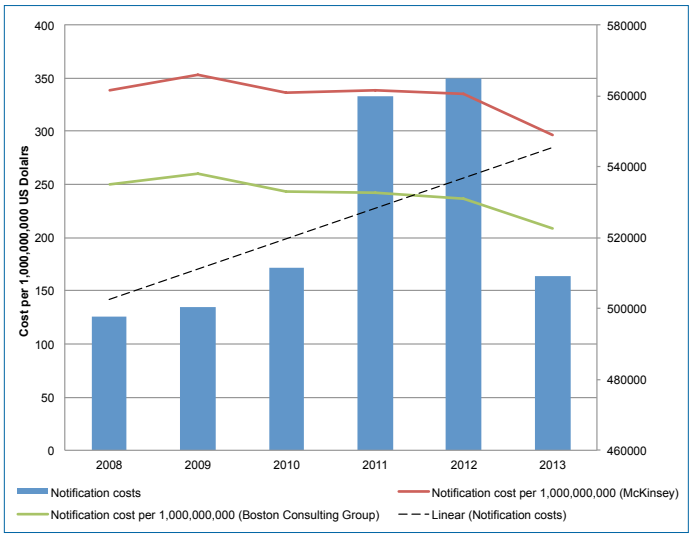

Figure 15, finally, presents data on the normalized and absolute trends in the costs that companies need to incur to inform individuals that their data has been breached. Here, despite the significant drop in the absolute cost of notification from $565,020 in 2012 to $509,237 in 2013, the general trend in the absolute numbers is toward higher and higher notification costs, as evidenced by the long-dash trend line in Figure 15. In contrast, the trend in both the normalized figures suggests that notification costs are actually declining between 2008 and 2013. In this case, the absolute numbers paint a picture of an increasingly costly security environment, while the normalized numbers suggest that the situation is actually getting better.

So, what conclusions can be drawn from these data on the cost of data breaches as a measure of the costs of cybercrime? Basically, the absolute numbers depict a worse cyber security situation than the normalized numbers. As with the measures for the vectors of cyber attack and the occurrence of cyber attacks, the absolute numbers create the perception that the security of cyberspace is worse than what is actually suggested by the more accurate normalized numbers.[16]

A few qualifiers are needed to temper these conclusions. The numbers in these cases are imperfect, as outlined above. Two points are worth reiterating. First, the economic contribution of the Internet to global GDP is likely larger than what is included in this study due to the assumption that the static, one-year estimates found in the McKinsey & Company and Boston Consulting Group studies are constant forward and backward throughout time. Secondly, the cost of data breaches is likely lower than what is found in these data, since the costs of cybercrime in the United States are, at least according to the Ponemon Institute's studies, consistently higher than the global average. Both of these qualifiers would actually strengthen the argument of this paper by lowering the various costs of cybercrime, while increasing the Internet's contribution to global GDP. Normalizing these lower numbers around this larger contribution suggests that the normalized trends would be even lower still.

In conclusion, in two of the six tests conducted in this section (post-breach response costs and notification costs), the absolute numbers point to a worsening situation, while the normalized numbers actually indicate that costs are declining. In three of the six cases (average cost per capita, overall organizational costs and lost business costs), both sets of numbers point to an improving situation, but the normalized numbers show the situation improving faster than the absolute numbers. Finally, in the last case (detection and escalation costs), both sets of numbers say the situation is getting worse, but the absolute numbers say that things are falling apart faster than the normalized numbers. Taken together, these findings once again indicate that the security of cyberspace (this time in terms of the costs of cybercrime) is actually better than the impression given by the commonly touted absolute numbers.

CONCLUSIONS AND POLICY RECOMMENDATIONS

What are the actual trends in cybercrime? Is the situation getting worse, better or staying roughly the same over time? We currently have a flawed picture of the security of cyberspace. Instead, a more accurate picture requires that the numbers on the occurrence of cybercrime be normalized around indicators that capture the growth and growing importance of cyberspace. To test this proposition, data on various indicators of the size of the cyberspace were collected, with a particular focus on users, points of interaction and the volume of online activity. Various measures of the occurrence of cybercrime were examined, with a focus on vectors of attack, the occurrence of attack and the cost of attacks. In every instance, the normalized numbers suggest that the security of cyberspace is better than what is found when one looks only at the absolute numbers. If you take lessons from the 13 normalizations, you find that six (almost half) point to a situation where the absolute numbers show a deteriorating situation while the normalized numbers actually show that things are getting better. In another six of the tests, both numbers show the situation as improving, but the normalized numbers

16 In the future, the absolute average cost of a data breach might steadily increase as more and more companies and state bureaucracies digitize their information. From a corporate or bureaucratic perspective, digitization promises many cost-saving and efficiency advantages. However, it also creates a larger potential cost if a data breach does occur. The future, in other words, might not be well predicted by the current trend of an improving cost scenario.

usually indicate that things are getter better sooner and faster compared to the absolute numbers. Finally, in the one case where both sets of numbers show the situation worsening, the absolute numbers still indicate that things are getting worse faster than the normalized numbers. Cyberspace, in other words, is more secure than is commonly perceived.

Research Conclusions

Any conclusions drawn from this research need to be qualified in light of the relatively poor data that is available for study. As pointed out above, an irony of cyber security research is that we live in an age of big data, but very little of this data on cyber security trends is actually publicly available. If the data underlying the study is inaccurate or subject to changes, then the conclusions themselves are also in need of revision. One likely scenario is that many of the indicators for cybercrime are probably higher than the data herein indicates. Software vulnerabilities go undisclosed. Cyber attacks go undetected. Data breaches go unreported. Nevertheless, this paper maintains that cybercrime in its three modalities (vectors, occurrence and costs) needs to be normalized in order to be properly understood, as has been done here. The numbers might be skewed, but they are definitely more accurate than the simple absolute figures.

Some interesting stories emerge when one looks more closely at some of the trends in the various figures. Obviously, the small number of data points restricts the confidence that we can have in any observations, but there are some suggestive tendencies. For instance, the data on botnets in Figure 7 shows that there has been a steady reduction in the number of botnets since 2008, both in absolute terms and as a proportion of the number of Internet users, email users and websites. This decline potentially suggests that people have become more conscious of the danger of having their computer commandeered for nefarious purposes and have taken steps (such as the use of anti-virus software or being more careful about sites visited) to prevent its occurrence. It could also suggest that there has been a more concerted and coordinated international effort by law enforcement agencies and private companies, such as Microsoft, to take down existing botnet networks and operators (Europol 2015). The cause of the decline is likely a mixture of both. Law enforcement efforts are knocking botnets offline, reducing the stock of infected computers, and individual actions may be slowing the rate of infection, reducing the growth of new botnets over time.

The absolute and normalized data in Figures 8 and 9 potentially tell an interesting story regarding whether cybercriminals or cyber security providers hold the initiative.[17] From 2009 to 2012, there is a rapid growth in both the absolute and the normalized number of web-based attacks, suggesting that cybercriminals are among the first to recognize the ways in which new technology can be exploited to make a profit. During 2012, the trend starts to reverse itself, and, in 2013 and 2014, both sets of numbers start to decline. This finding suggests two things. First, the Internet is growing rapidly and at a faster pace each year, which explains the rapid drop in the normalized number of attacks. Second, the decline in the absolute numbers also suggests that law enforcement efforts and individually undertaken security measures are effective at curbing the occurrence of web-based attacks. One interesting supposition that follows from this conclusion is that there are likely to be waves of web-based attacks in the future. Cybercriminals might quickly learn how to exploit new technologies, increasing crime, only to be followed by counteraction by individuals, businesses and law enforcement, which results in a decline in web-based assaults. This cyclical pattern, seen in a preliminary way in the data contained here, will likely be borne out at time goes on.

Lastly, as shown in Figures 10 and 11, both the average cost per breached record and the overall organizational cost of data breaches are declining in both absolute terms and normalized terms. Together, these two trends suggest that the number of data breaches overall might be declining, since both the average cost and the overall organizational cost are declining.[18] One limitation to what can be said on the basis of this data is that the available numbers exclude mega breaches, which compromise over 100,000 records in a single attack. It is also possible, therefore, that the costs of low-grade data breaches are declining because the size of your average data breach is increasing. At the same time, the available evidence suggests that most data breaches tend to be small and targeted at small-to-medium size enterprises (Gow, n.d.). In any event, based on the evidence presented here, the cost of data breaches seems to be decreasing.

Overall, these research results suggest that the security of cyberspace is actually better than what people might think from looking just at the absolute numbers. Assessing the precise effectiveness of cyber security measures given these trends is difficult because it requires a clear account of the counterfactual — that is, what would have happened in the absence of such policies. Put another way, an increasing trend might have actually increased even more or a declining tend might have been less pronounced had a particular policy not been in place. Despite this limitation, one conclusion that can be drawn from the

17 I am grateful to the reviewer for pointing out this interpretation to me.

18 I am again grateful to the reviewer for highlighting this interpretation of the data to me.

presented evidence is that current cyber security efforts are effective enough to limit the growth in vectors of attack, occurrence of attacks and the costs of attack to some extent. Since these signs of insecurity in cyberspace are not worsening too quickly in most cases, the rapidly growing size of cyberspace actually means that the overall security of cyberspace is, in a lot of cases, generally improving over time. In short, current cyber security policies, rather than being ineffective, are most likely actually helping the situation to a not insignificant degree.

Policy Recommendations

Several policy implications follow from the main finding of the paper. One cardinal mistake would be to assume that because the security of cyberspace seems to be improving, individuals, companies and governments do not need to act to protect themselves. If perceptions of the security of cyberspace are truly guided by the absolute figures (showing a poor and often worsening environment), then the real improvement in the security of cyberspace is probably driven in part by users' actions intended to counter this dangerous environment by increasing their IT security. More efforts along these lines are needed.

The following recommendations follow from the conclusions of this paper and the improvement of IT security more generally:

Focus on the individual. The weak point in most IT security systems is often the individual user and not the technical system itself. Spam and phishing emails are designed to capitalize on this weakness, which stems largely from a combination of a lack of knowledge and a likely moral hazard to do with individual responsibility for cyber security. In other words, many people do not know enough to not click the link in a phishing email and many people will likely click links in a work environment that they would never click at home because there is an IT staff to deal with the consequences and individual accountability for data breaches is inconsistent.

Detect and counter new vulnerabilities faster by relying on open source software where possible. Open source software, such as SSL, is often more secure than strictly proprietary programs because it can be examined by so many eyes, although some examples, such as the Heartbleed exploit in SSL, show that all software is vulnerable. Many, indeed most, individuals with a computer science or computer engineering background are committed to ideals of an open and free Internet. For that to occur, programs need to be secure. Available open source software tends to get examined more often because it is publicly available and this reveals vulnerabilities faster, leading to quicker fixes and more security. In comparison, proprietary programs tend, in general terms, to eventually get leaked, so criminals have access to that code too, but it is not examined by as many eyes, leading to less security on average (Clarke, Dorwin and Nash, n.d.).

Reduce the ability of state security agencies to retain zero-day exploits for law enforcement or national security purposes by requiring that they be disclosed to the software developer within a reasonable timeframe.[19] The US National Security Agency's (NSA's) policy toward zero-day vulnerabilities is one example of the problem of retention by state agencies. According to government sources, the NSA apparently must tell a company that it has discovered a zero-day exploit in its system (Zetter 2014). The major caveat to this requirement is that the NSA can closely guard its knowledge of the zero-day exploit if national security or law enforcement needs dictate (ibid.). Many, if not most, computer programs can be used the world over, so a zero-day exploit in nearly any program can theoretically have national security or law enforcement purposes because it could be used by adversaries of the United States. In the interregnum, while governments sit on zero-day exploits waiting for the chance to use them, the vulnerabilities can also be discovered by criminal elements and used to launch cyber attacks. Creating stricter rules around the disclosure of zero-day exploits, likely along the lines of a reasonable time frame for retention, perhaps on the order of six months to one year after discovery, would help limit the use of these exploits for criminal purposes.

Develop international agreements on spam, phishing emails and other forms of web-based attacks. Some agreements, particularly to do with spam, already exist. As the Internet spreads globally, the reach of these agreements must also spread. Bringing new nations into the potential agreements is also needed. In the case of some attacks, such as DDoS attacks, no agreement exists and there is much more to be done. Figuring out uniform rules to govern these different forms of cyber attack is an important step going forward.

Figure out ways — either through market mechanisms, state intervention or some combination of both — to spread out the costs of cybercrime. As shown above, the Internet contributes a lot more to the global economy than is taken away due to the costs of cybercrime. Overall average organizational costs due to data breaches, for example, is only a few thousand dollars for every billion dollars of global GDP that the Internet generates. At a global level, the costs of cybercrime are negligible, when you see how much the Internet is contributing to global GDP. Yet, these costs can cause individual firms considerable hardship. Cybercrime insurance is the likely way forward. In this vein, market mechanisms can help protect firms from the costs of cybercrime via a market-driven pricing mechanism that focuses on the risk and potential damage of a cyber attack. Governments could also intervene in the market to

19 I am grateful to Melissa E. Hathaway for suggesting this framing of this recommendation.

regulate the cost of cybercrime insurance and potentially even provide insurance themselves to help protect firms, possibly using a social, rather than a market, discount rate. In all likelihood, a combination of both market and state involvement in the insurance market is needed, especially in the short run, as the market is new and rife with imperfect information. The core idea is that some of the tremendous wealth generated by the Internet should be allocated toward insuring that the actual firms affected by data breaches are not completely destroyed by cybercrime.

Private companies whose operations rely on the Internet need to do more to protect themselves through training, capacity building and investment in IT security systems, at times supported by government grants in the case of small-to-medium sized enterprises (SMEs). The choice of who to target for a cybercrime is likely to be driven by two factors: the probability of successfully targeting the company and the size of the prize to be had.[20] Large companies tend to invest more in absolute terms in IT security than SMEs, making them more secure. At the same time, larger companies also offer a more tantalizing target than SMEs as they have more to steal. SMEs, in contrast, tend to invest less in IT security, making them easier targets, but are a less alluring prize for cybercriminals due to their smaller size. Essentially, all businesses are vulnerable. An important secondary implication is that rigorous efforts to provide for IT security at one level can actually displace criminals to another part of the economy, so if larger companies respond to insecurity in cyberspace with large investments in IT security, SMEs might be targeted more frequently. Recognizing this, there is a place for a government grant system to help SMEs develop better IT security so that they are not targeted disproportionately by cybercriminals.

Norton Symantec, Kaspersky Lab and other cyber security companies should start to collect and represent their data on cybercrime in normalized terms rather than as absolute or year-over-year figures. Understanding the level of insecurity that exists in cyberspace is vitally important and should form the basis of all public and corporate policy going forward. To get an accurate picture of the situation, the numbers on new vectors of attack, web-based attacks and the costs of cybercrime all need to be normalized around the growing size of cyberspace, otherwise a false impression is given, as shown in this paper. Norton Symantec, Kaspersky Lab and other companies of this sort could help provide valuable data for policy makers by developing — and publicly sharing— clear normalized numbers.

This paper has shown that the security of cyberspace is actually greater than the impression one gets when looking at the commonly used absolute figures. When the vectors of cyber attack, the occurrence of cyber attacks and the cost of data breaches are normalized around the growing size of cyberspace, the situation seems much less grim.

20 Another way to express this notion is that the probability of success (p = 0 to 1) discounts the value of what can be taken via a cyberattack (X = 0 through ∞). The basic cybercrime equation becomes P(X). For example, a cyberattack that is 50 percent likely to succeed and that is targeting a prize worth, say, 1,000,000 dollars results in 500,000 dollars' worth of prospective benefit (0.50[1,000,000] = 500,000). Likewise, a cybercrime that was 100 percent likely to succeed, but which the prize was only worth 500,000, would also be worth a total of 500,000 dollars to the cybercriminal. In short, the difficulty of the attack and the size of the prize both matter when a cybercriminal is picking a company to target.

Acknowledgements

This paper has benefitted from a number of capable eyes. Vivian Moser and Carol Bonnett of CIGI's publications team strengthened the language immensely. And, in no particular order, Andy Wyckoff, Simon Palamar, Laura DeNardis, Melissa Hathaway, Fen Osler Hampson, Gordon Smith, Bill Graham, David Clark and Leanna Ireland all provided terrific comments on the substance and style of the paper. Their efforts made it far stronger and sharpened the analysis and ideas. The remaining errors are mine and mine alone.

WORKS CITED

BBC News. 2015. "Cyber War Games to be Staged By UK and US." BBC News, January 16. www.bbc.com/news/uk-politics-30842669.

CIGI-IPSOS. 2014. "Global Survey on Internet Security and Trust." www.cigionline.org/internet-survey.

Cisco Systems. 2009. "Cisco Visual Networking Index: Forecast and Methodology, 2008-2013." www.cisco.com/web/BR/assets/docs/whitepaper_VNI_06_09.pdf.

———. 2010. "Cisco Visual Networking Index: Forecast and Methodology, 2009-2014." http://large.stanford.edu/courses/2010/ph240/abdul-kafi1/docs/white_paper_c11-481360.pdf

Clarke, Russel, David Dorwin and Rob Nash. n.d. "Is Open Source Software More Secure?" http://courses.cs.washington.edu/courses/csep590/05au/whitepaper_turnin/oss%2810%29.pdf.

Dean, David, Sebastian Digrande, Dominic Field, Andreas Lundmark, James O'Day, John Pineda and Paul Zwillenberg. *2012. The Connected World: The Internet Economy in the G20: The 4.2 Trillion Growth Opportunity.* Boston Consulting Group. www.bcg.com/documents/file100409.pdf.

Europol. 2015. "Botnet Taken Down through International Law Enforcement Cooperation." Europol, February 25.

Evans, Dave. 2011. "The Internet of Things: How the Next Evolution of the Internet is Changing Everything." Cisco White Paper. April. www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.

Finkle, Jim, Soham Chatterjee and Lehar Maan. 2014. "EBay Asks 145 Million Users to Change Passwords after Cyber Attack." Reuters, May 21. www.reuters.com/article/2014/05/21/us-ebay-password-idUSBREA4K0B420140521.

Gandalf Group. 2014. "The 36th Quarterly C-Suite Survey: Cyber-Security, Trade Agreements and Foreign Investment." www.gandalfgroup.ca/downloads/2014/C-Suite%20Presentation%20Q3%202014%20Oct%2020%20TC.pdf.

Griffith-Jones, Stephany, Eric Helleiner and Ngaire Woods. 2010. *The Financial Stability Board: An Effective Fourth Pillar of Global Economic Governance.* Special Report. Waterloo, ON: CIGI.

Gow, Brad. n.d. "Data Security Breaches: More Reach and Frequency Requires More Diligence." Zurich. www.zurich.com/NR/rdonlyres/C4FC10D0-2156-42F8-84E7-63C3BF69B6B6/0/Tech_Cold2_DataBreach.pdf.

Internet Live Stats. 2015a. "Internet Users." www.internetlivestats.com/internet-users/.

———. 2015b. "Total Number of Website." www.internetlivestats.com/total-number-of-websites/#trend.

ITU. 2015. Statistics. www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx.

Kaspersky Lab. 2008. "Kaspersky Security Bulletin 2008." http://securelist.com/analysis/kaspersky-security-bulletin/36241/kaspersky-security-bulletin-statistics-2008/.

———. 2009. "Kaspersky Security Bulletin 2009." http://securelist.com/analysis/kaspersky-security-bulletin/36284/kaspersky-security-bulletin-2009-statistics-2009/.

———. 2010. "Kaspersky Security Bulletin 2010." http://securelist.com/analysis/kaspersky-security-bulletin/36345/kaspersky-security-bulletin-2010-statistics-2010/.

———. 2011. "Kaspersky Security Bulletin 2011." http://securelist.com/analysis/kaspersky-security-bulletin/36344/kaspersky-security-bulletin-statistics-2011/.

———. 2012. "Kaspersky Security Bulletin 2012." http://securelist.com/analysis/kaspersky-security-bulletin/36703/kaspersky-security-bulletin-2012-the-overall-statistics-for-2012/.

———. 2013. "Kaspersky Security Bulletin 2013." http://media.kaspersky.com/pdf/KSB_2013_EN.pdf.

———. 2014. "Kaspersky Security Bulletin 2014." http://cdn.securelist.com/files/2014/12/Kaspersky-Security-Bulletin-2014-EN.pdf.

Norton Symantec. 2009. *Internet Security Threat Reports*. http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiv_04-2009.en-us.pdf.

———. 2010. *Internet Security Threat Reports*. http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xv_04-2010.en-us.pdf.

———. 2011. *Internet Security Threat Reports*. www4.symantec.com/mktginfo/downloads/21182883_GA_REPORT_ISTR_Main-Report_04-11_HI-RES.pdf .

———. 2012. *Internet Security Threat Reports*. www.trustico.com/news/internet_security_reports/internet_security_report_2012_04.en-us.pdf.

———. 2013. *Internet Security Threat Reports*. www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf.

———. 2014. *Internet Security Threat Reports*. www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf.

———. 2015. *Internet Security Threat Reports*. www4.symantec.com/mktginfo/whitepaper/ISTR/21347932_GA-internet-security-threat-report-volume-20-2015-social_v2.pdf.

O'Toole, James. 2014. "Mobile Apps Overtake Desktop Internet Usage in U.S." CNN Money, February 28. http://money.cnn.com/2014/02/28/technology/mobile/mobile-apps-internet/.

Pélissié du Rausas, Matthieu, James Manyika, Eric Hazan, Jacques Bughin, Michael Chui and Rémi Said. 2011. *Internet Matters: The Net's Sweeping Impact on Growth, Jobs and Prosperity.* McKinsey & Company. May. www.mckinsey.com/insights/high_tech_telecoms_internet/internet_matters.

Ponemon Institute. 2011. *2011 Cost of Data Breach Study: United States.* Ponemon Institute Research Report. www.ponemon.org/local/upload/file/2011_US_CODB_FINAL_5.pdf.

———. 2013. *2013 Cost of Data Breach Study: A Global Analysis.* Ponemon Institute Research Report. www4.symantec.com/mktginfo/whitepaper/053013_GL_NA_WP_Ponemon-2013-Cost-of-a-Data-Breach-Report_daiNA_cta72382.pdf.

———. 2014. *2014 Cost of Data Breach Study: Global Analysis.* Ponemon Institute Research Report. www-935.ibm.

com/services/us/en/it-services/security-services/cost-of-data-breach/.

Radicati Group. 2013. "Email Market 2013-2017." www.radicati.com/wp/wp-content/uploads/2013/11/Email-Market-2013-2017-Executive-Summary.pdf.

Richwine, Lisa. 2014. "Sony's Hacking Scandal Could Cost the Company $100 Million." *Business Insider*, December 9. www.businessinsider.com/sonys-hacking-scandal-could-cost-the-company-100-million-2014-12.

Royal Pingdom. 2009. "Internet 2008 in Numbers." http://royal.pingdom.com/2009/01/22/internet-2008-in-numbers/.

———. 2010. "Internet 2009 in Numbers." http://royal.pingdom.com/2010/01/22/internet-2009-in-numbers/ .

———. 2011. "Internet 2010 in Numbers." http://royal.pingdom.com/2011/01/12/internet-2010-in-numbers/.

———. 2012. "Internet 2011 in Numbers." http://royal.pingdom.com/2012/01/17/internet-2011-in-numbers/.

———. 2013. "Internet 2012 in Numbers." http://royal.pingdom.com/2013/01/16/internet-2012-in-numbers/.

Silver-Greenberg, Jessica, Matthew Goldstein and Nicole Perlroth. 2014. "JPMorgan Chase Hacking Affects 76 Million Households." *The New York Times*, October 2. http://dealbook.nytimes.com/2014/10/02/jpmorgan-discovers-further-cyber-security-issues/?_r=0.

Statista. 2015. "Number of Smartphones Sold to End Users Worldwide from 2007 to 2014 (in million units)." www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/.

Statistics Brain. 2015. "Google Annual Search Statistics." www.statisticbrain.com/google-searches/.

StatsCounter. 2015. "Top Desktop, Console and Tablet Browers per Country, Oct 2014." http://gs.statcounter.com/#browser-ww-monthly-201410-201410-map.

Verisign. 2009. *Domain Name Industry Brief* 6 (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb09.pdf.

———. 2010. *Domain Name Industry Brief* 7 (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb10.pdf.

———. 2011. *Domain Name Industry Brief* 8 (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb-2011.pdf.

———. 2012. *Domain Name Industry Brief* 9 (1). Verisign. www.verisigninc.com/assets/domain-name-brief-march2012.pdf.

———. 2013. *Domain Name Industry Brief* 10 (1). Verisign. www.verisigninc.com/assets/domain-name-brief-april2013.pdf.

———. 2014. *Domain Name Industry Brief* 11 (1). Verisign. www.verisigninc.com/assets/domain-name-report-april2014.pdf.

———. 2015. *Domain Name Industry Brief* 12 (1). Verisign. www.verisigninc.com/assets/domain-name-report-march2015.pdf.

Woodcock, Bill and Vijay Adhikari. 2011. "Survey of Characteristics of Internet Carrier Interconnection Agreements." Packet Clearing House. May 2. www.pch.net/resources/papers/peering-survey/PCH-Peering-Survey-2011.pdf.

Zetter, Kim. 2014. "U.S. Gov Insists it Doesn't Stockpile Zero-Day Exploits to Hack Enemies." *Wired*, November 17. www.wired.com/2014/11/michael-daniel-no-zero-day-stockpile/.

CIGI PUBLICATIONS
ADVANCING POLICY IDEAS AND DEBATE

Global Commission on Internet Governance

The Global Commission on Internet Governance (GCIG) was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. The two-year project conducts and supports independent research on Internet-related dimensions of global public policy, culminating in an official commission report that will articulate concrete policy recommendations for the future of Internet governance. These recommendations will address concerns about the stability, interoperability, security and resilience of the Internet ecosystem. Launched by two independent global think tanks, the Centre for International Governance Innovation and Chatham House, the GCIG will help educate the wider public on the most effective ways to promote Internet access, while simultaneously championing the principles of freedom of expression and the free flow of ideas over the Internet.

Finding Common Ground
A Briefing Book Prepared for the Global Commission on Internet Governance

This briefing book contextualizes the current debate on the many challenges involved in Internet governance. These include: managing systemic risk — norms of state conduct, cybercrime and surveillance, as well as infrastructure protection and risk management; interconnection and economic development; and ensuring rights online — such as technological neutrality for human rights, privacy, the right to be forgotten and the right to Internet access.

The Regime Complex for Managing Global Cyber Activities
GCIG Paper Series No. 1
Joseph S. Nye, Jr.

Tipping the Scale: An Analysis of Global Swing States in the Internet Governance Debate
GCIG Paper Series No. 2
Tim Maurer and Robert Morgus

Legal Mechanisms for Governing the Transition of Key Domain Name Functions to the Global Multi-stakeholder Community
GCIG Paper Series No. 3
Aaron Shull, Paul Twomey and Christopher S. Yoo

Legal Interoperability as a Tool for Combatting Fragmentation
GCIG Paper Series No. 4
Rolf H. Weber

Innovations in Global Governance: Toward a Distributed Internet Governance Ecosystem
GCIG Paper Series No. 5
Stefaan G. Verhulst, Beth S. Noveck, Jillian Raines and Antony Declercq

The Impact of the Dark Web on Internet Governance and Cyber Security
GCIG Paper Series No. 6
Tobby Simon and Michael Chertoff

On the Nature of the Internet
GCIG Paper Series No. 7
Leslie Daigle

Understanding Digital Intelligence and the Norms That Might Govern It
GCIG Paper Series No. 8
David Omand

ICANN: Bridging the Trust Gap
GCIG Paper Series No. 9
Emily Taylor

A Primer on Globally Harmonizing Internet Jurisdiction and Regulations
GCIG Paper Series No. 10
Michael Chertoff and Paul Rosenzweig

Connected Choices: How the Internet is Challenging Sovereign Decisions
GCIG Paper Series No. 11
Melissa E. Hathaway

Solving the International Internet Policy Coordination Problem
GCIG Paper Series No. 12
Nick Ashton-Hart

Net Neutrality: Reflections on the Current Debate
GCIG Paper Series No. 13
Pablo Bello and Juan Jung

Addressing the Impact of Data Location Regulation in Financial Services
GCIG Paper Series No. 14
James M. Kaplan and Kayvaun Rowshankish

Cyber Security and Cyber Resilience in East Africa
GCIG Paper Series No. 15
Iginio Gagliardone and Nanjira Sambuli

ABOUT CIGI

The Centre for International Governance Innovation is an independent, non-partisan think tank on international governance. Led by experienced practitioners and distinguished academics, CIGI supports research, forms networks, advances policy debate and generates ideas for multilateral governance improvements. Conducting an active agenda of research, events and publications, CIGI's interdisciplinary work includes collaboration with policy, business and academic communities around the world.

CIGI's current research programs focus on three themes: the global economy; global security & politics; and international law.

CIGI was founded in 2001 by Jim Balsillie, then co-CEO of Research In Motion (BlackBerry), and collaborates with and gratefully acknowledges support from a number of strategic partners, in particular the Government of Canada and the Government of Ontario.

Le CIGI a été fondé en 2001 par Jim Balsillie, qui était alors co-chef de la direction de Research In Motion (BlackBerry). Il collabore avec de nombreux partenaires stratégiques et exprime sa reconnaissance du soutien reçu de ceux-ci, notamment de l'appui reçu du gouvernement du Canada et de celui du gouvernement de l'Ontario.

For more information, please visit www.cigionline.org.

ABOUT CHATHAM HOUSE

Chatham House, the Royal Institute of International Affairs, is based in London. Chatham House's mission is to be a world-leading source of independent analysis, informed debate and influential ideas on how to build a prosperous and secure world for all. The institute: engages governments, the private sector, civil society and its members in open debates and confidential discussions about significant developments in international affairs; produces independent and rigorous analysis of critical global, regional and country-specific challenges and opportunities; and offers new ideas to decision-makers and -shapers on how these could best be tackled from the near- to the long-term. For more information, please visit: www.chathamhouse.org.

CIGI MASTHEAD

Executive

President	Rohinton P. Medhora
Director of the International Law Research Program	Oonagh Fitzgerald
Director of the Global Security & Politics Program	Fen Osler Hampson
Director of Human Resources	Susan Hirst
Vice President of Public Affairs	Fred Kuntz
Director of the Global Economy Program	Domenico Lombardi
Vice President of Finance	Mark Menard
Chief of Staff and General Counsel	Aaron Shull

Publications

Managing Editor, Publications	Carol Bonnett
Publications Editor	Jennifer Goyder
Publications Editor	Vivian Moser
Publications Editor	Patricia Holmes
Publications Editor	Nicole Langlois
Graphic Designer	Melodie Wakefield
Graphic Designer	Sara Moore

Communications

Communications Manager	Tammy Bender	tbender@cigionline.org (1 519 885 2444 x 7356)

CIGI
67 Erb Street West
Waterloo, Ontario N2L 6C2
tel +1 519 885 2444 fax +1 519 885 5450
www.cigionline.org

CHATHAM HOUSE
The Royal Institute of International Affairs
10 St James's Square
London, England SW1Y 4LE, United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org