The Role of Internet Service Providers in Botnet Mitigation
An Empirical Analysis Based on Spam Data[1]

Michel van Eeten[a)], Johannes M. Bauer[b),] Hadi Asghari[a)],
Shirin Tabatabaie[a)], Dave Rand[c)]

a) Delft University of Technology, Faculty of Technology, Policy and Management, the Netherlands
{m.j.g.vaneeten h.asghari s.tabatabaie}@ tudelft.nl

b) Michigan State University, Department of Telecommunication, Information Studies, and Media, USA
bauerj @ msu.edu

c) Trend Micro Incorporated, USA
Dave_Rand @ trendmicro.com

Botnets – networks of machines infected with malicious software – are widely regarded as a critical security threat. Measures that directly address the owners of the infected machine end users are useful, but have proven insufficient to reduce the overall problem. Recent studies have shifted attention to key intermediaries – most notably, Internet Service Providers (ISPs) – as control points for botnet activity. Surprisingly little empirical information is available to assess the claim that ISPs are an important control point, as well as related claims, for example, that large ISPs are worse cybercitizens than smaller ones. This paper is a first effort to go beyond generalized arguments by dissecting the diversity of ISPs and the number of infected machines in their networks. As most of the current spam is sent through botnets, the origin of spam messages provides us with a proxy for detecting infected machines. Using a global dataset of 138 million unique IP addresses that connected to a spam trap in the period 2005-2008, we have analyzed in detail the geographic patterns, time trends, and differences at the level of countries and ISPs. This data underlines the key position of ISPs as intermediaries. For example, in our dataset just 10 ISPs account for around 30 percent of all unique IP addresses sending spam worldwide; 50 ISPs account for over half of all sources. For the first time, the patterns in infected machines are connected to other data, such as the size of the ISPs and the country in which they are located. Using bivariate and multivariate statistical approaches we investigate empirically the effects of country-level policy measures on the number of unique IP addresses sending spam at the ISP level. The data reveals wide differences between ISPs in the relative number of infected machines, sometimes up to three orders of magnitude. Whereas the overall number of infected machines is largely driven by the size of the user base, we also find limited evidence that public policies to improve cybersecurity have the desired mitigating effects. Our findings confirm some of the claims made in the research literature but refute others.

[1] The authors would like to acknowledge the financial support of the Organisation for Economic Co-operation and Development (OECD). They also would like to thank three anonymous reviewers for their helpful feedback and comments which have contributed to clarifying the arguments presented in the paper.

Background

The internet economy is highly dependent on information and network security. Estimates of the direct damage caused by internet security incidents vary wildly, but typically range in the tens of billions of US dollars per year for the U.S. alone (e.g., US GAO 2007; Bauer et al. 2008). In addition, all stakeholders in the information and communication system incur indirect costs of possibly even larger magnitude, including costs of prevention. While this damage is related to a wide variety of threats, the rise of malicious software ('malware') and botnets are seen as a, if not *the*, most urgent security threat we currently face.

If recent estimates are correct, around 5 percent of all machines connected to the Internet may be infected with malware (BBC News 2007; House of Lords 2007; Moore et al. 2009). The fact that the owners of these machines often do not know their machines are compromised is part of the problem. Malware may be distributed and used in many ways, including email messages, USB devices, infected websites, malicious advertising, and browser vulnerabilities (Jakobsson and Zulfikar 2008).

The massive number of compromised machines has allowed the emergence of so-called 'botnets' – networks of thousands or even millions of infected machines that are remotely controlled by a 'botnet herder' and used as a platform for attacks as well as fraudulent and criminal business models, such as the sending of spam and malicious code, the hosting of phishing sites, to commit click fraud, and the theft of confidential information.

While originating in criminal behavior, the magnitude and impact of the malware threat is also influenced by the decisions and behavior of legitimate market players such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. As security comes at a cost, tolerating some level of insecurity is economically rational. Market players make their decision based on the perceived costs and benefits of a course of action. In many markets these benefits also reasonably reflect the resource costs and benefits of a course of action to society at large. However, economic research and policy analysis have identified situations in which this correspondence is weakened and systematic gaps between private and social costs and benefits of security exist, a situation for which the term "externalities" is used.

Botnet mitigation by Internet Service Providers

Recent research suggests that infected end user machines, in particular those of home users and small and medium-size enterprise (SME) users, are a key source of security externalities (Van Eeten and Bauer 2008). In contrast to larger corporate users, these groups often do not select desirable levels of protection.

Measures that address end users directly – including awareness raising and information campaigns – are useful, but they have proven to be insufficient to reduce the overall problem. Recent studies have therefore shifted attention to key intermediaries, most

notably, ISPs – in the sense of access providers, not providers of hosting or other services. As access providers to end users, they form, to some extent, a natural control point for the effects of infected machines. Anderson et al. (2008, pp. 50-54) argue that liability for infected machines should be assigned to the ISPs, rather than to the consumers who own the machines. The authors also propose to impose statutory damages on ISPs that do not respond promptly to requests for the removal of compromised machines.

Of course, the fact that ISPs can potentially mitigate this threat, does not mean that they *should* mitigate it. They are not the source of the externality and they have to bear substantial direct and indirect costs if they do internalize the externalities of their customers. Nevertheless, in a variety of countries, ISPs are now explicitly assuming some responsibility for botnet mitigation. Industry collaborative efforts like the Internet Engineering Taskforce (IETF) and the Messaging Anti-Abuse Working Group (MAAWG) have prepared sets of best practices for the remediation of bots in ISP networks. Under pressure from the government, Australia's largest ISPs are preparing a voluntary code of conduct that includes contacting infected customers and filtering their connection.

Within the OECD, other countries have indicated they are pursuing similar lines of action. A related initiative in Germany is the establishment of a government-funded call center to which ISPs can direct customers in need of support to disinfect their machines. The largest ISPs in the Netherlands – with an aggegate market share of over 90 percent – have entered into a covenant that expresses their commitment to mitigate botnet activity in their own networks. They claim that their organizations already have practices in place where they contact and in some cases quarantine customers whose machines are infected with malware. While this may be true, there is currently no data available that indicates the scale on which these practices are being carried out.

Scale is critical, however. There are indications that ISPs only deal with a fraction of the infected machines in their networks. For example, in an earlier study we found that a large ISP with over four million customers contacted around 1,000 customers per month (Van Eeten and Bauer 2008, p. 29). Typical estimates of security researchers put the number of infected machines at around five percent of all connected machines at any point in time (Moore et al. 2009, p. 5). This would translate into about 200,000 infected machines for this specific ISP. Even if we reduce the estimated infection rate to one percent, that still implies 40,000 infected machines. This stands in stark contrast to the 1,000 customers that the ISP claimed to be contacting – even when we optimistically assume that all contacted customers either willing and able to clean up their infected machine or are being quarantined.

To reiterate: We are not claiming that ISPs should contact all the owners of infected machines. That is a matter for policy development to consider, taking into account the costs and benefits of mitigation, for ISPs, their customers, as well as society at large. We are simply stating that there is an urgent need to collect data, beyond the generic claims of ISPs that they are contacting customers and quarantining infected machines. This data

should inform us not only about the extent to which ISPs can mitigate, are actually mitigating, and how they perform relative to each other.

To this end, the paper sets out to empirically answer the following questions: First, to what extent are ISPs critical control points for botnet mitigation? Second, to what extent do they perform differently relative to each other, in terms of the number of infected machines in their networks? Third, and last, to what extent can we explain the differences in performance from the characteristics of the ISPs or the environment in which they are located?

Before turning to these questions, we first outline the research approach, as well as its limitations. At the heart of the research is data from a spam trap that has logged around 138 million unique IP addresses of machines that connected to it. The raw data was parsed to associate IP addresses with ISPs and countries. We then examine the intermediairy position of ISPs. Surprisingly, in our dataset, just 50 ISPs account for half of all unique IP addresses of infected machines worldwide. We also explore the differences among ISPs in the extent in which their networks harbor infected machines. These differences are substantial, even when corrected for the size of the customer base of the ISP. To explain these differences, we employ bivariate and multivariate statistical approaches. Among others, using ISPs as the unit of analysis, we investigate empirically the effects of country-level policy measures on the number of unique IP addresses sending spam. We conclude with a discussion of the implications of our findings for current efforts to mobilize ISPs in botnet mitigation.

Research approach

There is no authoritative data source identifying infected machines around the world. Roughly, there are two types of sources: (1) data collected external to the botnet, identifying infected machines by their behavior, such as sending spam or participating in distributed denial of service attacks; (2) data collected internal to the botnet, identifying infected machines by intercepting communications within the botnet itself, for example by infiltrating the command and control infrastructure of the botnet.

Each known source has its own strengths and weaknesses. The first type typically uses techniques such as honey pots, intrusion detection systems and spam traps. It has the advantage of identifying machines across a wide range of botnets. The drawback is that there are issues with false positives and negatives. The second type typically intercepts botnet communications by techniques such as redirecting traffic or infiltrating IRC channel communication. The advantage of this approach is accuracy: bots connecting to the command and control server are really infected with the specific type of malware that underlies the botnet. The downside is that measurement only captures infected machines within a single botnet. Given the fact that the number of botnets is estimated to be in the hundreds (Zhuang et al. 2008), such data may not be representative of the overal population of infected machines.

This study draws upon data from spam traffic – a source of the first type. The originating IP address of spam messages provides us with a useful source of proxy data for infected machines (see also Zhuang et al. 2008). The bulk of all spam messages are sent through botnets. Estimates published during the period under study put the figure at around 80 to 90 percent of the total amount of spam (Ironport 2006; Messagelabs 2009). The originating IP address of a spam message is therefore very likely to indicate the presence of at least one infected machine.

Our data is drawn from a spam trap – an Internet domain set up specifically to capture spam, whose email addresses have never been published or used to send or receive legimitate email traffic. In the period of 2005-2008, the trap has received 63 billion spam messages and incoming SMTP connections from about 138 million unique IP adresses worldwide.

Of course, not all spam comes from infected machines and not all infected machines send spam. The first issue points to the risk of false positives. As mentioned above, 80 to 90 percent of all incoming spam originates from a botnet. We have reason to believe that for the spam received by our trap this ratio is even higher. The trap is located at a small and relatively old generic top-level domain. Tactics to distribute spam through other means than botnets, such as "snowshoe spamming", are typically more targeted and use fresher addresses, in part because these tactics are more costly than the use of botnets. In other words, this spam would not be captured by our trap and not lead to false positives. More importantly, at a later stage of the analysis, we split all spam sources in two categories, depending on whether the network in which the source is located belongs to an ISP or not. We focus our analysis on the first category, which eliminates a lot of potential false positives, namely spam from sources such as webmail providers, hosting providers and university networks. In short, we have reason to assume that the impact of false positives is limited. The second issue – not all infected machines send spam – points to the risk of false negatives, of undercounting infected machines. Our data undoubtedly suffers from undercounting, as do all existing data sources. That being said, sources external to botnets, such as spam traps, are less affected by this limitation than internal data sources, because they identify infected machines across a wide range of botnets. In that sense, these sources can be considered the most representative of the overall population.

For each unique IP address that was logged by the spam trap, we looked up the Autonomous System Number (ASN) and the country where it was located, using the MaxMind geoIP database. As both ASN and geoIP information change over time, we used historical records to establish the orgin for the specific moment in time at which the message was received. We also recorded the number of spam messages sent from each source.[2] This effort resulted in two time series of variables: unique IP addresses and spam

[2] The IP address of the incoming SMTP connection attempts were checked against a blacklist of known spam sources. In case the address was on the list, the connection was refused. To conservatively estimate how many messages these refused connections would have contributed to the spam volume, we calculated the daily average number of message sent per accepted connection attempt. Given that refused connections were from known spam sources, the number of messages these sources would have sent if the connection was accepted is likely to be higher than the daily average.

volume, both per ASN and per country. The former is more directly related to the number of infected machines. The latter variable is useful to balance some of the shortcomings of the former – a point to which we return in a moment.

We have conducted extensive triangulation efforts to compare our data to the publicly available reports of security and anti-spam service providers. Most of the public data relates to the relative spam volume of countries. The commercial reports present different numbers, sometimes substantially different numbers. The patterns and distributions that we found were within the range reported by the commercial providers.

We then set out to identify the ISPs to which the ASNs belonged. To the best of our knowledge, there is no existing database that maps ASNs onto ISPs. This is not surprising. Estimates of the number of ISPs vary from around 4,000 – based on the number of ASNs that provide transit services – to as many as 100,000 companies that self-identify as ISPs – many of whom are virtual ISPs or resellers of capacity of other ISPs.

So we adopted a variety of strategies to connect ASNs to ISPs. First, we used historical market data on ISPs – wireline, wireless and broadband – from TeleGeography's GlobalComms database. We extracted the data on all ISPs in the database listed as operating in a set of 40 countries, namely all 30 members of the Organisation for Economic Co-operation and Development (OECD), plus five "accession candidates" and five so-called "enhanced-engagement" countries. This resulted in data on 200 ISPs (see Appendix 1).

The process of mapping ASNs to ISPs was done manually. First, using the GeoIP data, we could identify which ASNs where located in each of the 40 countries. ASNs with one percent of their IP addresses mapped to one of the 40 countries were included in our analysis. Next, we listed all ASNs in a country that were above a threshold of 0.5 percent of total spam volume for that country.

We then checked the ASNs on this list against the list of ISPs in that country, as per the TeleGeography database. We used historical WHOIS records for each ASN to lookup its name and then consulted a variety of sources to see which, if any, of the TeleGeography operators it matches. In many cases, the mapping was straightforward. In other cases, more information was needed – for example, in case of ASNs of ISPs that had since been acquired by another ISP. In those cases, we mapped the ASN to its current owner.[3]

While we believe this to be a robust approach to answer our empirical questions, it has certain limitations – most notably, the effects of Network Address Translation (NAT), dynamic IP addresses with short lease times and port 25 blocking. The question is how these practices affect the number of machines that are represented by a unique IP address.

[3] We mapped ASNs by going down the list of top spam-sending ASNs in each country, ranked by volume, until one of the following conditions was met: (1) 95 percent of the spam originating from that country had been covered; or (2) the number of ASNs covered is five times the number of ISPs in that country, as listed in the TeleGeography database; or (3) the next ASN contributes less than 1 percent of spam originating from that country and less than 0.01 percent of spam worldwide.

NAT means sharing a single IP address among a number of machines. This potentially underrepresents the number of infected machines, as they all show up as a single address. Dynamic IP addresses with short lease times implies that a single machine will have multiple IP addresses over time. This overrepresents the number of infected machines. Both of these practices counteract each other, to some extent. This limits the bias each of them introduces in the data, but this does not happen in a consistent way across different networks. Earlier research by Stone-Gross et al. (2009) has demonstrated that in different countries, there are different ratios of infected machines to unique IP addresses – the so-called "churn rates".

We have two ways to robustly control for the potential bias that these churn rates introduce in our data. First, we look at the volume of spam in addition to the number of unique sources. If there are many machines behind a single IP address, the spam volume should be relatively high, even if it looks like a single source. If there is one machine behind many IP addresses, the spam volume should be relatively low. We have calculated the ratio of spam volume to unique sources in our data. The Spearman correlation between these ratios and the churn rates reported by Stone-Gross et al. (2009) is very high, namely -0.88. This suggests that spam volume can control for churn. A second way to control for it is by calculating the daily averages of the number of unique sources for ISPs. Research by Moore et al. (2002) found that, because of DCHP churn, IP addresses are not an accurate measure the number of infected machines on timescales longer than 24 hours. We therefore ran all our analysis also using the daily averages and found that all patterns discussed below are consistent with daily averages.

For all the analyses we discuss in this paper, we have always checked whether the pattern we found also persisted when using both of these controls. For the sake of brevity, we focus our discussion on the number of unique sources. When spam volume or the daily averages show a different pattern, we explicitly include it in the discussion. Where they are not mentioned, they are consistent with the findings as reported here.

A final limitation is the use of port 25 blocking by ISPs. The effect of port blocking is that infected machines can no longer directly send email to the wider internet, but have to go through the ISP's outgoing email servers. This affects both the number of sources as well as the spam volume. The ISP's network may harbor housands of infected machines, but they can no longer reach the spam trap directly and thus do not reveal their IP address through spam distribution. There is one important way in which the attackers themselves compensate for this problem: when the bots notice they cannot connect anymore via port 25, they start to send spam via the ISP's official outgoing email servers. In various cases where port blocking was introduced, we saw that it led to a brief reduction of outgoing spam, only to return to the previous spam volume within about a month. It is difficult for the ISP to prevent this from happening, as each bot sends out a relatively low level of spam, and thus rate limits and similar controls do not pick up on it. The effect of this tactic is that here, too, spam volume provides the abilty to cross check our findings, to some extent. In other cases, port blocking is an unavoidable limitation to our data. If the spam volume remains consistently lower, port blocking obscures the presence of infected machines. That being said, the effect of the bias is not wholly unreasonable. The ISPs that

adopt port blocking improve their ranking in terms of botnet activity compared to those that don't – which is not without merit, given that the measure of port blocking is part of many guidelines on best security practices for ISPs.

Are ISPs critical control points?

The most important reason to focus on ISPs as intermediaries is that they are viewed as critical control points. To some extent, this is obvious. The ISP's customers can only send and receive traffic via the ISP, which creates a natural bottleneck to mitigate malicious activity of the customers' machines. However, there are two important assumptions that are rarely explicitly acknowledged. First, to what extent are infected machines actually located within the networks of ISPs? In other words, what about the machines in use by, for example, hosting providers, application service providers, webmail providers, university networks and corporate networks? If ISPs can only control a minor portion of the infected machines, it undermines the argument to focus on them, more than the other players, as the key intermediaries in the fight against botnets.

The second assumption behind the idea to focus on ISPs as control points is that the burden will be put on the relevant ISPs. We are most familiar with the legitimate ISPs, well-known brands that together possess the bulk of the market share. These organizations are identifiable, reachable and stable enough to be brought into some form of collaborative process or under a regulatory regime. However, as security incidents have often pointed out, there is also a class of so-called "grey" and "rogue" ISPs. This class may have a disproportionate part in the impact of botnet activity. They also typically evade, intentionally or not, the normal processes through which collective action is brought about. If we stimulate ISPs to do botnet mitigation, voluntarily or through some type of policy measures, the burden will not fall onto this class of ISPs. In other words, treating ISPs as control points implicitly assumes that the problem exists for the most part within the networks of the legimate providers that have most of the market share; not in the margins of the market, which is teeming with large numbers of small ISPs that are often shortlived and difficult to survey, let alone reach through public regulation or self-regulation.

As far as we know, these assumptions have never been emprically tested. Our data allows us to do just that. As explained above, we are working with a set of 200 ISPs in the wider OECD – 30 member countries and 10 associated countries. This set consists of the ISPs that collectively posses the bulk of the market share in these countries. We first looked at the portion of the total number of unique sources of spam that can be attributed to these ISPs. Over the period of 2005-2008, between 63-69 percent of all global sources were located within networks of the 200 ISPs. For spam volume, the numbers are slightly lower: 50-64 percent (see Figure 1). If we look at the total number of sources in the 40 countries where the ISPs are located, that ratio is, of course, even higher: 77-82 percent. This confirms the first assumption, namely that the bulk of infected machines are located in the networks of the larger, predominantly retail ISPs – rather than hosting providers, corporate networks, application service providers. These appears to be little, if any, volatility in this pattern.

Figure 1: Percentage of sources compared to global total

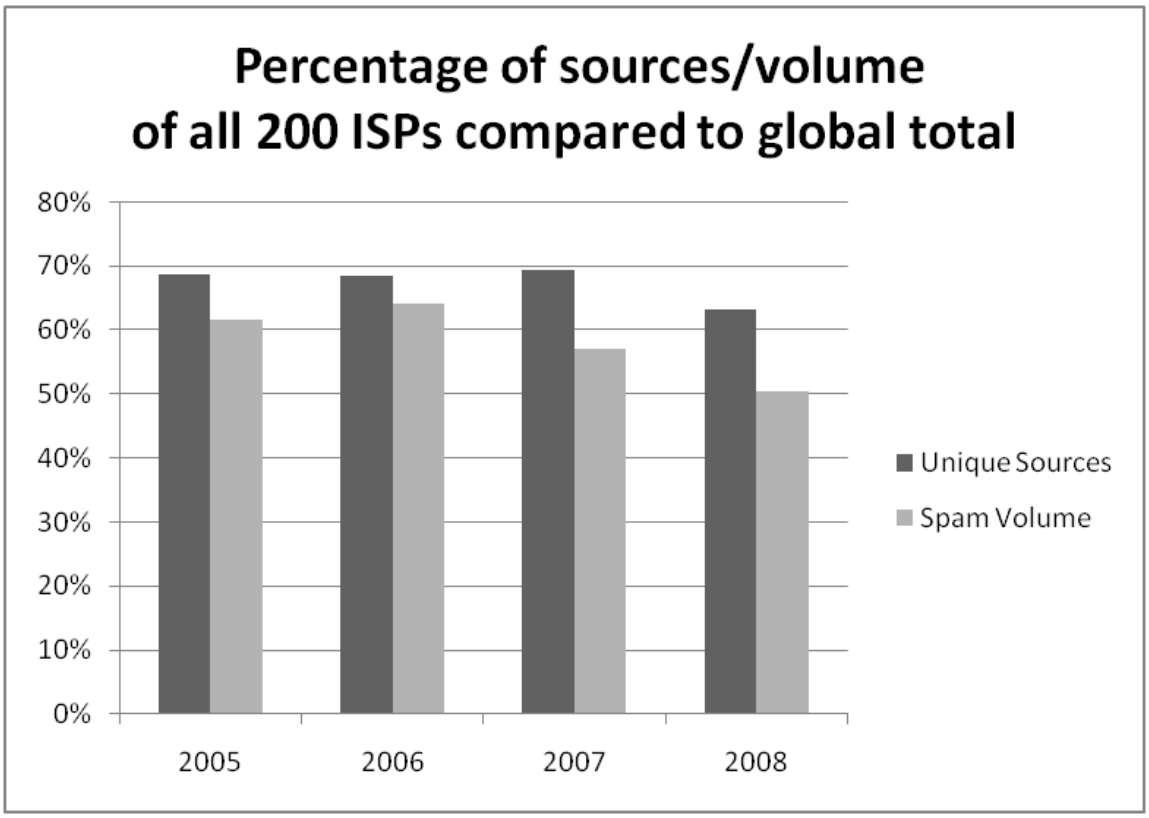

It is interesting to note that these ratios vary significantly across countries. On the high end, we have countries like Israel, Turkey and Italy, where in 2008 over 96 percent of all sources are located with ISPs. On the low end, we find Canada, with around 42 percent, which may be explained by the fact that Canada has a large hosting provider industry.

It is also interesting to look at the distribution of sources within this set of 200 ISPs. That gives us a sense of the validity of the second assumption. If we rank the ISPs in order of the number of unique sources in their networks in 2008, we find that the 10 highest ranking ISPs account for around 30 percent of all unique sources worldwide (figure 2). The top 50 ISPs account for over half of all sources worldwide. In light of the fact that there are 30,000 ASNs and anywhere between 4,000-100,000 ISPs, this is a remarkable finding. We also see that the curve flattens quickly. Adding the next 150 ISPs captures only an additional 8 percentage points of sources worldwide. This confirms the second assumption.

In light of the thousands of players that are involved, collective action would seem an almost futile pursuit, given all the typical problems of free rider behavior and weakest-link security. For botnet mitigation, however, the task of combating infected machines turns out to have more manageable proportions, institutionally speaking. Our findings strongly suggest that the more established and visible ISPs are indeed the ones who form

Figure 2: Percentage of sources located in top ISPs

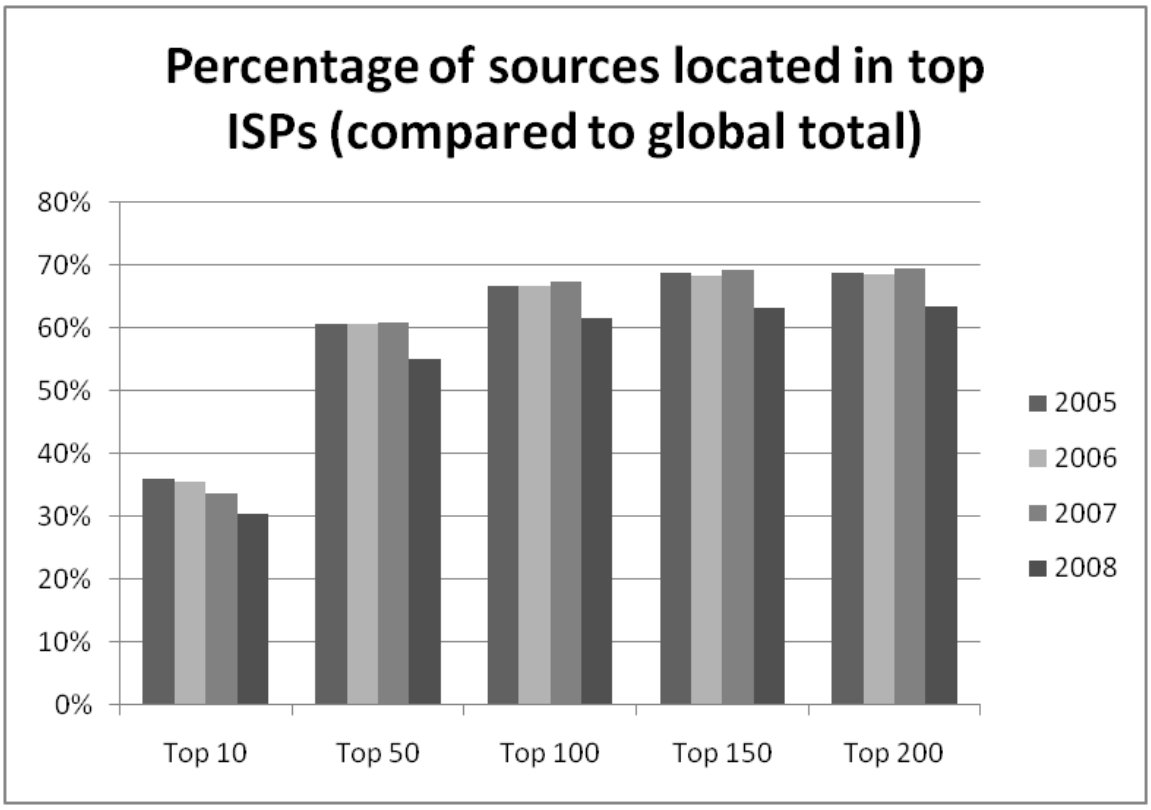

critical control points, not the thousands of smaller players that would be difficult, if not impossible, to reach through collaborative or regulatory efforts.

Of course, none of this is to say that improving botnet mitigation has suddenly become an easy task. Nor are we arguing that the same pattern holds across other types of malicious activity. Many types of criminal activity do, in fact, thrive because of weakest-link problems among ISPs – as business models such as bulletproof hosting have demonstrated.

Do ISPs perform differently in terms of botnet mitigation?

A lot has been written about the incentives of ISPs, or lack thereof, to improve security (House of Lords 2007; Van Eeten and Bauer 2008; Bauer and Van Eeten 2009; Moore et al. 2009). Various incentives have been identified, some enhancing security, others working against it. It is not at all clear what the net effect is of these incentives on ISP's behavior, nor whether this effect varies significantly across ISPs. Another way to frame this problem is to ask how much discretion ISPs have in mitigating botnets. If they are subject to similar incentives but have little organizational freedom to respond to them, then we would expect similar performance in this area. However, if they do have discretion and can respond differently, diverse performance outcomes will be observable.

Figure 3: Unique sources and number of subscribers of ISPs in the OECD (2008)

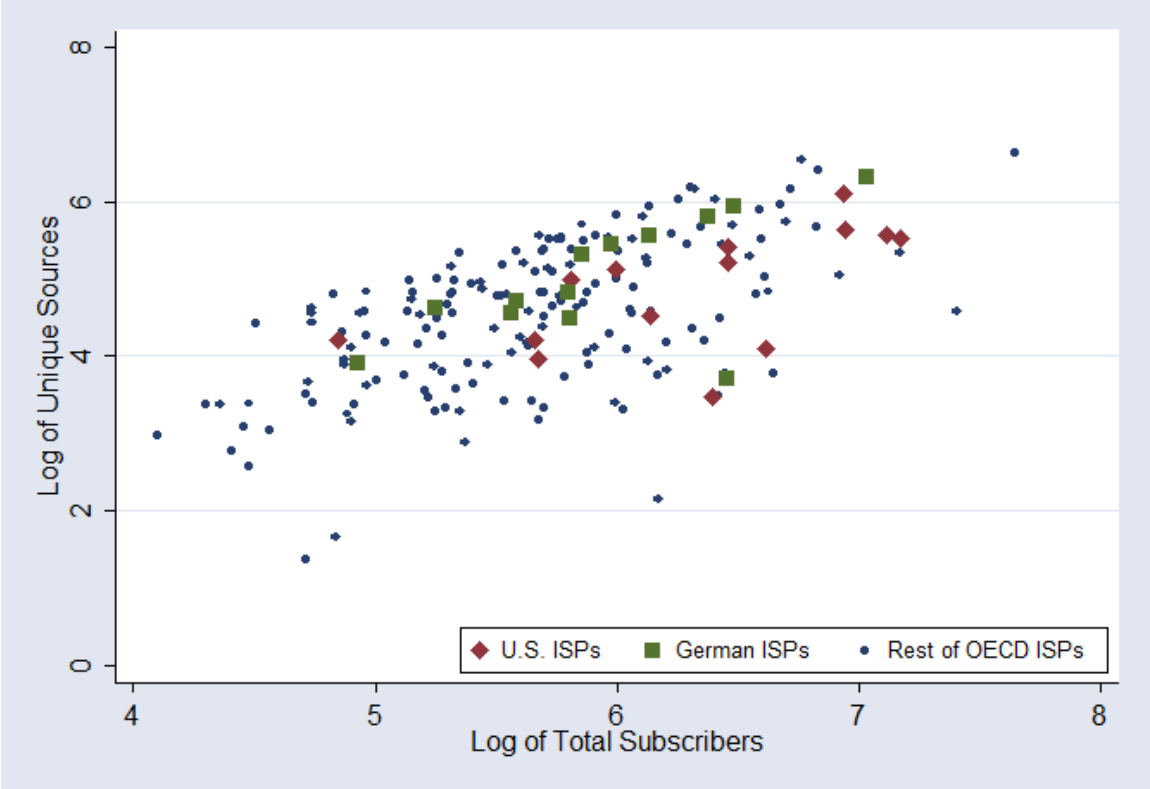

The first factor to take into account is size. Our dataset includes a range of ISPs of varying size. In 2005, the smallest ISP had 3,000 customers and the largest 21 million customers, with the median at approximately 252,000. In 2008 the numbers were higher, with the smallest ISP showing 12,600 and the largest 44.3 million customers. The median in 2008 was at 500,750 customers. For some analysis, the dataset was split into small and large ISPs, with key statistics of the latter by about an order of magnitude higher than for the former group. Obviously, other things being equal, ISPs with more customers will experience more infections. If we look at ISP performance – as measured in number of unique sources and spam volume – and rank the ISPs according to their size, this becomes immediately visible (see figure 3 for the 2008 findings). We can see a nearly linear relationship, when both variables are transformed logarithmically (R^2 is 0.43 for the period 2005-2008). That being said, there is still considerable variability.

Across the board, we see a difference of two orders of magnitude, sometimes even higher, in the number of infected machines within networks of ISPs of similar sizes. This is not a matter of outliers. The coefficient of variation – basically the ratio of the standard deviation to the mean – is well above 1, both for the number of unique sources and spam volume. Other incentives may be country-specific, such as the cost of legal requirements or the cost of customer support. But even within countries we see substantial differences in performance. In the U.S. and Germany, for example, we still see at least one order of magnitude difference, often more, among ISPs of similar size (figure 3).

A third aspect is the performance of ISPs over time. Although we did not yet perform a detailed empirical analysis of the factors that explain differences in the dynamic response of ISPs to infections on their networks, we conducted an aggregate analysis of ISPs in the set of worst offenders. The Venn diagrams in figure 4 illustrate overlaps in the 50 worst performing ISPs between 2005 and 2008 both for the absolute and relative number of sources. There is some variation in membership in the total set and the various sub-sets. For example, a total of 66 ISPs were in the top 50 in one of the four years based on the number of infected machines and 77 ISPs were in the top 50 based on the number of infected machines per subscriber. However, we also observe a stable core of 34 ISPs that had the highest number of infected machines on their network during all four years (9 ISPs were in the set in three years, 12 in two years, and 11 in only one year). With regard to infected machines per subscriber, 28 ISPs were in the set during all four years, 15 were in it during three years, 9 during two, and 25 during only one year.

Figure 4: Variability in top 50 sources of spam 2005-2008

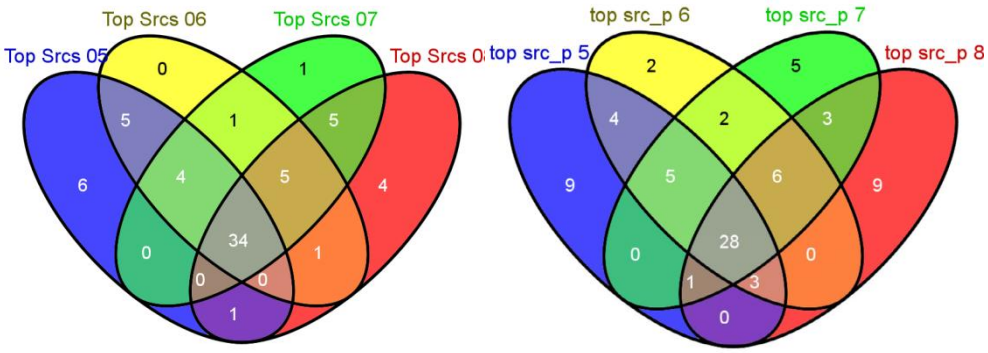

(a) Variability in absolute number of sources

(b) Variability in relative number of sources

The size distribution of ISPs in the two core sets of poor performers throughout the entire time period is compatible with the overall picture gained from the statistical analyses. The smallest ISP with regard to absolute number of sources had about 480,000 subscribers and the largest 44.3 million. With regard to sources per subscriber, the smallest ISP reported about 28,600 subscribers and the largest one 5,8 million subscribers. Within these sets, although there is considerable variation, larger ISPs on average did better than smaller ISPs. The patters remains the same when taking average daily numbers of sources, either in absolute or relative, per subscriber terms.

All of this suggests that ISPs have significant discretion to decide how they engage in botnet mitigation and that their organizational incentives lead to different choices, even when working under a common set of institutional incentives, such as defined by the legal framework of a country. This point is reinforced when we look at the differences between countries, rather than ISPs. At the country level, our data measures the total spam output of ISPs and non-ISPs. As players with very different records within one

country are aggregated, country performance data show less variance than individual organization data. Consequently, at that level of analysis, the number of internet users explains around 70 percent of the variance in performance. As ISPs do perform very differently under comparable institutional incentives and economic circumstances, this suggests that country-level mitigation measures, while necessary, will not be sufficient unless they also address the organizational incentives and realign both. In the next section, we explore the extent to which we can explain these differences among ISPs.

Explaining the differences among ISPs

Advanced information and communication technologies form a highly interrelated ecosystem. Like other actors, ISPs respond to economic and non-economic incentives. Most generally speaking, incentives are the factors that individual and organizational decision-makers take into account. Given the highly dynamic nature of this ecosystem, the observations reported in the previous section could be the complex outcomes of the varied responses by ISPs to the problems of botnets without an underlying stable pattern. However, if the phenomenon had certain regularities this knowledge could be utilized to improve cybersecurity. We therefore formulated a simplified conceptual model of the ecosystem around ISPs and subjected it to empirical analysis. Figure 4 represents a stylized model of the factors that influence botnet activity: the security measures adopted by an ISP, the level and virility of cybercriminal attacks, technological factors, and user behavior. Other factors, such as the behavior of software vendors and registrars, also impact this ecosystem, but they are outside the scope of this study (see van Eeten and Bauer 2008 for an in-depth discussion). An ISP's decisions to adopt security measures are influenced by factors related to the institutional and organizational environments. These groups of factors are linked in multiple feedbacks so that they co-evolve over time. For example, stronger security efforts by an ISP may reduce botnet activity but also result in stronger efforts by cybercriminals to find new attack vectors. As our units of analysis are ISPs, it is important to take the national context into account. However, cybercrime is a transborder phenomenon and the international context is therefore also relevant.

The incentive structure of a particular ISP is shaped by institutional and organizational factors. These two sets of factors are interrelated in many ways. For example, a regulation obliging an ISP to undertake certain security measures has cost implications at the organizational level. Likewise, the failure of ISPs to adopt a sufficient level of security-enhancing measures increases the likelihood that institutional responses might be sought. It is nonetheless useful to distinguish them, as institutional incentives can be designed by policy makers whereas organizational ones are typically shaped by managers (often in response to institutional incentives). Overall, the resulting incentive structure under which an ISP operates consists of a mix of contradictory forces, some increasing efforts to mitigate botnets (other things being equal) and others weakening them (other things being equal). For example, if higher botnet activity increased the risk of being blacklisted this constitutes a positive incentive– i.e., an incentive to improve security and to mitigate botnet activity. In contrast, the cost of acting against infected machines is a negative incentive, as higher costs reduce botnet mitigation efforts. Depending on the

strength of the relation between an incentive and the effort to mitigate botnet, incentives fall on a continuum from high-powered (strong) to low-powered (weak).

The level of effort that ISPs exert on botnet mitigation depends on the relevant set and the relative strength of positive and negative incentives. Relevant institutional factors include the legal and regulatory framework in which ISPs operate, the market structure and the associated competitive pressures, and the conditions in related markets, e.g. for security technology. Relevant organizational factors include the size of the customer base, the organization of the abuse desk, and the cost of various security measures. Which incentives will be perceived as relevant by an ISP is influenced by its business model. Commercial ISPs will primarily respond to incentives that have direct and indirect implications for their bottom line. Likewise, rogue ISPs deriving most of their business from activities related to cyberfraud and cybercrime will also primarily respond to economic incentives (Van Eeten and Bauer 2008). In both cases, non-economic incentives, such as peer pressure and peer recognition, may play a role. These types of incentives are often seen to be subordinate to economic incentives. This need not be the case, however. When peer pressure takes the form of blacklisting, it has economic effects that can be quite significant, such as rising cost of customer support, when customers experience the effects of blacklisting and start calling their ISP. The relative weights of relevant incentives could be different for non-profit ISPs or cooperatives but even such ISPs do not have unlimited resources and will have to pay attention to economic factors. All ISPs will therefore be influenced by the incentives identified in Figure 4, which interact to jointly influence an ISP's botnet mitigation effort.

The signs in Figure 4 refer to the direction of the incentive, other things being equal. A positive sign indicates that the incentive has likely a positive effect on the level of botnet mitigation by an ISP. The strength of an incentive is quite a different issue and may depend on the presence of complementary incentives. For example, laws providing a base for action against spammers will only be effective if they are also enforced actively. Likewise, the effectiveness of liability rules, which are often mentioned as a possible course of action, will depend on whether or not the required burden of proof can be met. Because such enforcement encounters great difficulties, leading legal scholars tend to be skeptical whether liability rules are workable (e.g., Spindler 2007). The effectiveness of incentives and the interaction between them will also be influenced by the national context. Of particular importance are the diffusion of broadband service, the income level of a country, the education level of the population, and the diligence of law enforcement.

Agents in this ecosystem usually have an incomplete view of the relevant facts and/or of the consequences of particular actions and make their choices within these informational boundaries ("bounded rationality"). Moreover, while there will be some shared ("common") information, part of the incomplete information will be asymmetrically distributed among the stakeholders. Agents even in otherwise similar organizations may therefore respond differently to the same set of institutional incentives if their knowledge differs. Therefore, one would expect a diverse set of responses to the institutional and organizational incentives under which ISPs operate. Despite this diversity of responses, the effect of incentives can nevertheless be systematically examined.

Figure 4: Conceptual framework

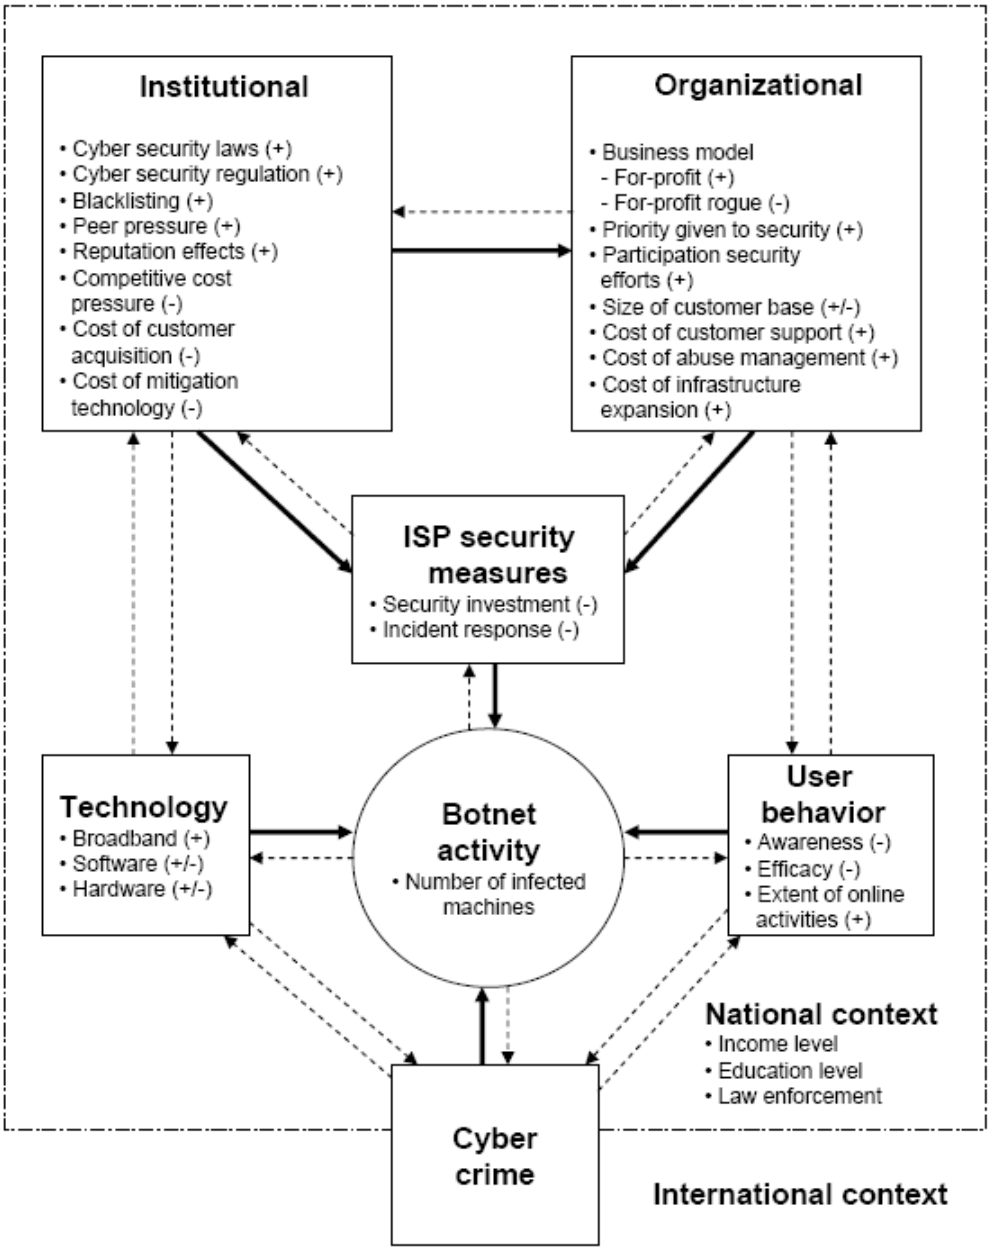

Due to data constraints, for purposes of empirical analysis, the conceptual framework discussed in the previous paragraphs was further simplified to a more manageable empirical model. No data on the number of infected machines was available. We approximated it with two measures: the number of unique IP addresses emitting spam and the total number of spam messages originating from an ISP during a specific time period. Drawing on the conceptual framework discussed above, a large number of

Figure 5: Empirical framework

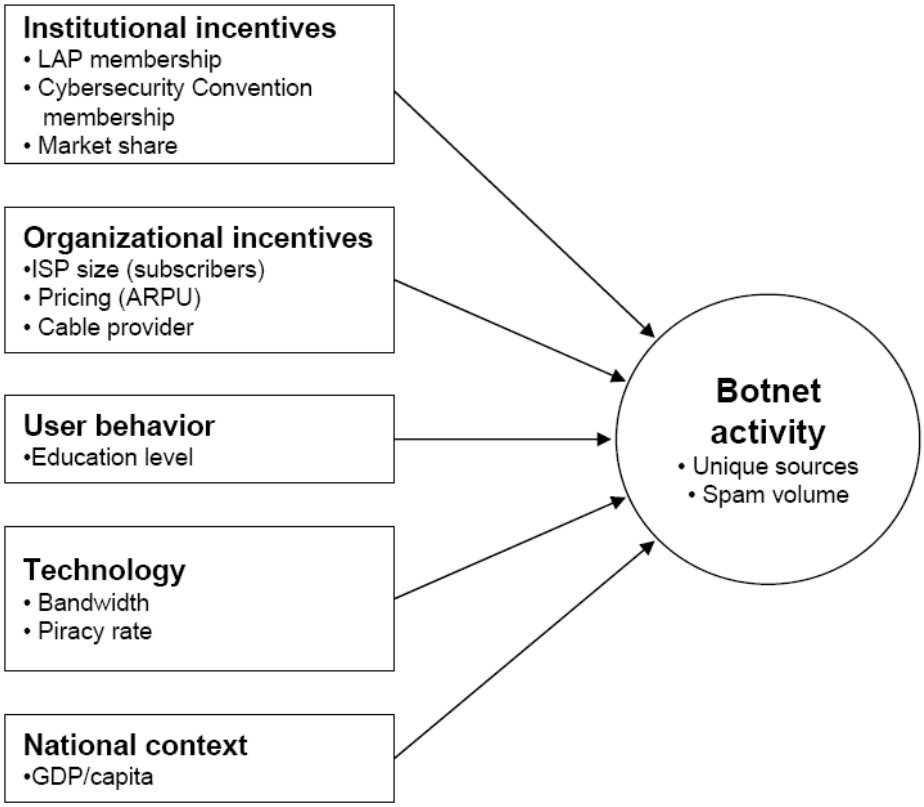

variables that could either serve as direct measures of as proxies for the independent variable were used. In this paper, only variables that were seen as relevant based on the conceptual approach and that yielded a statistical contribution are reported. In addition, control variables were introduced to take factors related to technology, user behavior, and the national context into account. As cybercrime is a globally mobile phenomenon, we proceeded on the assumption that all ISPs are targeted at a comparable rate (although the level of botnet activity is influenced by an ISP's security efforts as well as other control factors and will thus vary). The empirical model is displayed in Figure 5.

Data for the independent and control variables was collected from several sources, including the World Bank's World Development Index database, the UN Human Development Reports, the Software Business Alliance, and TeleGeography's GlobalComms database (see Appendix 2). Where possible, data was triangulated against other sources, such as the International Telecommunication Union's World Telecommunications Indicator database. In addition to the 63 billion spam messages from 138 million unique sources, which were parsed, aggregated, and attributed to ISPs and countries in the way discussed above, we were able to assemble a panel of annual observations for 2005-2008 for 40 countries. Although we were able to gather considerable evidence, it was not possible to generate empirical data for all the variables suggested by the theoretical model forcing us to work with proxies where available.

However, in some cases, such as prices for internet access services, the empirical model was constrained by lack of data.

Empirical findings

The dynamic nature of botnets raises many methodological challenges. Three approaches, each with its own advantages and disadvantages, were used to test hypotheses derived from the theoretical framework: (1) bivariate methods, (2) multivariate methods using pooled data, and (3) panel data analysis. In each case, we set out to explain relative differences in botnet activity among the ISPs. Unless noted otherwise, we used as the dependent variables the number of infections per subscriber – i.e., unique sources per subscriber and spam volume per subscriber.[4]

Bivariate analysis offers a first crude look onto the relations between independent variables and the proxies for botnet activity. However, it has to be kept in mind that bivariate statistics neglect the influence of other factors that may play a role and therefore may attribute too much influence to a single independent variable. They therefore grant only a preliminary understanding of the data structures. The robustness of findings needs to be checked against multivariate statistical methods. The results of this analysis are presented in Table 1.

That ISPs (as opposed to other types of players, such as hosting providers or corporations operating a network with its ASN) play a central role in botnet activity was already discussed. Likewise, the great variability among ISPs was already discussed. In addition to these findings, our data indicate the following (see Asghari 2010 for a more detailed discussion):

- There is a widely held belief that larger ISPs show worse security performance, as they face much less peer pressure. For instance, Moore, Clayton, and Anderson (2009, p. 10) state that "...very large ISPs are effectively exempt from peer pressure as others cannot afford to cut them off. Much of the world's bad traffic comes from the networks of these 'too big to block' providers." In contrast to this belief, our dataset indicates that, while larger ISPs emit more spam in absolute numbers, relative to size their performance is on average slightly better than that of smaller ISPs.
- Another claim is that lower average revenue per user (ARPU) is a sign of higher financial pressure that might result in less attention to security. Our data suggests that ARPU and relative security performance are unrelated.
- Given differences in networking technology and user base, one might hypothesize that cable service providers can enhance their security performance easier than DSL providers. Our data indicates an eight percent better lower incidence of unique sources for cable companies. The volume of spam, however, is similar for both types of providers, which might reflect that cable subscriptions have higher average bandwidths than DSL subscriptions.
- Bivariate analysis indicates that countries that have joined the London Action Plan (LAP) have, on average, 12 percent fewer bot infections. Likewise, being a signatory

[4] See Appendix 3 for the descriptive statistics of the variables and Appendix 4 for the pair wise correlations between the independent variables.

Table 1: Bivariate test results for number of unique sources per subscriber*

Subject	Independent variables	Statistical instrument	N	Results *pooled, sources*	Results *pooled, volume*	Results
Effects of ISP size	total_sub	Spearman's rho	*N=741*	*sig_1=0.000* *ρ = -0.170*	*sig_1=0.000* *ρ = -0.182*	**Negative relation**
	market_share			*sig_2= .820*	*sig_2= .179*	**No relation**
Effects of ARPU	rev_persub	Spearman's rho	*N=194*	*sig = 0.275*	*sig=0.770*	**No relation**
Cable vs. DSL providers	srv_cable	t-test	*N=665*	*sig = 0.000* *diff = .0766*	*sig = 0.506*	**Cable providers have fewer sources**
Effects of regulation	lap_mem	t-test,	*N=741*	*sig_1 = 0.000* *diff = .120*	*sig_1 = 0.000* *diff = 33.8*	**LAP members have fewer sources**
	cyberconv_mem	Kruskal-Wallis, t-test		*sig_2 = 0.000* *diff = .129*	*sig_2 = 0.000* *diff = 33.6*	**CC members have fewer sources**
Effects of piracy	piracy_rate	Spearman's rho	*N=740*	*sig = 0.000* *ρ = 0.391*	*sig = 0.000* *ρ = 0.342*	**Positive relation**
Effects of bandwidth	int_bpp	Spearman's rho	*N=386*	*sig = 0.000* *ρ = -0.232*	*sig = 0.421*	**No relation**
Effects of user education	educ_ix	Spearman's rho	*N=547*	*sig = 0.000* *ρ = -0.381*	*sig = 0.000* *ρ = -0.261*	**Negative relation**

* The bivariate statistical tools used are comparison of means and measures of association. In this table, *sig* is the statistical significance of the test result (values under 0.05 are considered statistically significant results); *ρ* is the rank correlation coefficient, and is a measure of the degree of association of two variables (between -1 to 1); *diff* is the difference between the averages of the two sample groups. Due to lack of normality in the dependent variable, often non-parametric tests were employed.

to the Council of Europe's Convention on Cybercrime is negatively correlated with botnet infections. Neither of these initiatives targets botnets directly. However, one could argue that membership is a proxy for the overall commitment of a country's government to enhance cybersecurity – and thus of a broader set of measures undertaken. Earlier research by Wang and Kim (2009) provided some evidence in support of this effect, though they presume a somewhat tenuous direct causal link between the Convention and cybercrime, rather than interpreting membership of the Convention as a proxy variable. However, factors correlated with a country's willingness to sign these agreements could also be at work both for the Convention as well as the LAP.

- A frequently stated claim is that countries with higher rates of software piracy also have higher botnet activity. At the bivariate level, our data supports that a moderate positive relation exists between piracy and botnet activity.
- Bandwidth is often seen as enabler of malware (e.g., OECD 2009). However, our data does not support that claim at the bivariate level and we did not find an indication that increased use of broadband connections does "automatically" translate into a higher number of bot infections – measured either in the number unique sources or spam volume.

Table 2: Pooled regression results for unique sources of spam (log transformed)*

Dependent variable: unq_srcs_log	All ISPs (β)	Small ISPs (β)	Large ISPs (β)
totsub_log	0.650 **	0.558 **	0.428 **
market_sh	0.077 *	-0.020	0.143 **
srv_cable	-0.086 **	-0.192 **	-0.010
cyber_mem	-0.068	-0.023	-0.072
lap_mem	-0.139 **	-0.125 *	-0.234 **
educ_ix	-0.024	-0.111	0.028
piracy_rate	0.104 **	0.102	0.165 *
_cons	0.466	0.527	-0.498
N	639	300	339
R^2_{adj}	51.4%	36.6%	34.0%

* ISPs were split into the two subsets based on having a total number of subscribers below or above 395,000. The chosen cut-off point is the median number of subscribers in the pooled set of observations. Reported betas (β) are standardized (except _cons). Significance levels: * ≤ 0.05, ** ≤ 0.01.

- Lastly, we were interested in whether higher education levels are associated with lower levels of botnet activity. In the bivariate analysis, a negative effect of higher education on botnet activity is indeed visible.

To overcome the limitations of bivariate analyses, multiple regression analyses were conducted. With four years of information available, the data could be examined from different perspective (although only a few selected findings are reported here), including cross sectional analyses of annual data, pooled data, and panel data estimation. In a pooled data design, the driving methodological assumption is that the same generative process explains all observations, independent of the ISP and/or the year. This implies that parameters do not vary between the units of analysis. Although this is a strong assumption, in the present case, where all ISPs are subject to a relentless stream of attacks of a predominantly global nature, it is not entirely unrealistic. A recent study found that half of the detected botnets included machines in over 30 countries. Some botnets even control machines in over 100 countries (Zhuang et al. 2008).

The relative measures of botnet activity (number of unique sources per subscriber or spam volume per subscriber) are more intuitive, because we want to compare ISPs. But the downside of using dependent variables normalized with the size of the ISP is that they require us to use instrument variables so as to not violate a key assumptions of the linear regression model. For this reason, transformations of the variables are used. Moreover, to gain insights into the factors driving the total number of infected machines, we first specified a model using the absolute number of unique sources for each ISP – transformed by using a log function.[5] The double-log specification has the advantage that β-coefficients can be interpreted as elasticities.

[5] We used a logarithmic transformation because the order of magnitude of the number of subscribers is more important the absolute number – i.e., we would expect security practices to differ between an ISP

The findings from the pooled regression analysis are presented in the second column of Table 2. The model explains about 51 percent of the variation among ISPs in the number of unique sources (49 percent when using spam volume as dependent variable). The model is largely congruent with the bivariate findings, except for the impact of the Cybercrime Convention and education, which now are found to be weaker and non-significant.

To explore the relationship between ISP size and botnet activity that was found in the bivariate analysis, we divided the set of ISPs into two groups: small and large ISPs. The results are presented in the third and fourth column of Table 2. The initial finding is confirmed. The elasticity of unique sources to changes in the number of subscribers is higher in smaller ISPs – that is, a one percent increase in the number of subscribers leads to a higher increase in the percentage of infected sources for small ISPs (0.56) than large ISPs (0.43). Simply put, smaller ISPs are, on average, doing slightly worse.

The effect that cables providers have lower infection rates appears to hold primarily for smaller ISPs. If indeed, as we hypothesized earlier, this effect is tied to automation that could explain why we do not see the effect of cable versus DSL for large ISPs, as they are more likely to already have automation in place because of their size. The effect of regulatory activity – as measured by the proxy of LAP membership of the country in which the ISP is located – is stronger in large ISPs. This fits with the earlier observation that large ISPs are more within the reach of governmental efforts to improve cybersecurity.

The next step in multivariate analysis was to model the relative performance of ISPs, i.e., the amount of botnet activity corrected for size of the ISP. [6] The results are presented in Table 3. This model explains about 36 percent of the variance, notwithstanding the many factors at play in the botnet phenomenon. However, it also clearly indicates that other factors are at work, pointing to the highly dynamic nature of the phenomenon – not in the last place because of volatile patterns caused by the attackers. An analysis of the error terms indicated the presence of heteroscedasticity, which weakens but does not invalidate the findings. It is a possible indication that other factors that are not yet included in the model, may be at work.

As incentives typically do not work in isolation from each other, we also introduced interaction terms to capture the joint effects of selected factors. Interaction terms appear among the country level variables, indicating that they change in 'configurations', as is often the case with institutional and demographic variables.

---

with 50,000 subscribers versus one with 500,000, but not between ISPs with 5 million and 5.5 million subscribers.

[6] This variable was transformed using a square root function, because of the law of diminishing returns at work for this variable – e.g., completely infection-free ISP networks are non-existent, but as the number of infections goes up, it becomes increasingly difficult to add additional infections – i.e., it is all but impossible to achieve a 100 percent infection rate.

Table 3: Pooled regression results for unique sources per subscriber (sqrt transformed)

```
      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F( 11,   652) =   34.68
       Model |  12.2717863    11  1.11561693           Prob > F      =  0.0000
    Residual |  20.9735399   652  .032168006           R-squared     =  0.3691
-------------+------------------------------           Adj R-squared =  0.3585
       Total |  33.2453262   663   .05014378           Root MSE      =  .17935

------------------------------------------------------------------------------
  src_per_sq |      Coef.   Std. Err.      t    P>|t|                     Beta
-------------+----------------------------------------------------------------
  totsub_log |  -.0336822   .0122981    -2.74   0.006                -.1072685
   srv_cable |  -.4511269   .1281703    -3.52   0.000                -.9433911
  icblXsubln |     .06528   .0229551     2.84   0.005                 .7614554
   cyber_mem |  -.0403519   .0213033    -1.89   0.059                -.0801629
     lap_mem |   5.055042   1.300163     3.89   0.000                 11.16164
 piracy_rate |  -.0395952   .0137778    -2.87   0.004                -3.022274
     educ_ix |  -3.083531   .9688927    -3.18   0.002                -.8433481
    ilapXedu |  -5.597067   1.355777    -4.13   0.000                -11.91356
    ilapXpir |  -.1252716   .0205351    -6.10   0.000                -11.96453
    ieduXpir |   .0438092   .0144469     3.03   0.003                 2.782703
ilapXeduXpir |   .1393075   .0218696     6.37   0.000                 12.18444
       _cons |   3.535415    .938184     3.77   0.000                        .
------------------------------------------------------------------------------
```

Although some of the findings differ from the simple bivariate analysis, there is also considerable congruence. With the exception of membership in the Cybercrime Convention, all variables in Table 3 are significant at the one percent level. In this model specification, the parameters of the total number of subscribers of the ISP, cable service provision status, membership in the Cybercrime Convention, and education levels of users all were negative, indicating that these factors mitigated botnet activity. In the multivariate setting, the parameter sign of the piracy rate, however, switched to negative and the parameter sign of LAP membership to positive. As these factors may interact with others, we tested several specifications of interaction effects. Of these, interactions of LAP membership with education and piracy generated negative parameter signs (indicating, for example, a botnet mitigating effect of the interaction of LAP membership and education but, less convincingly, also of LAP membership and piracy rate). This implies that some of the findings are sensitive to the specification of the model and therefore less robust than other findings that do not change.

The third approach used panel data methods. Panel data allow taking advantage of the cross-sectional and time-series dimensions of the data. In other words the method takes advantage of the fact that data originated from different ISPs and at different points in time. In our case, we used a fixed effects model, which relaxes the assumption of the pooled data approach that one generative process drives the botnet phenomenon and allows ISP-specific differences. With only four years of observations, though, panel data estimation has inherent limitations. Moreover, variables that do not change within one country during the four years (e.g., institutional incentives such as LAP or Cybercrime Convention membership) cannot be used and are therefore dropped in the estimation procedure. Of the three methods, panel data estimation therefore is the most challenging

Table 4: Panel regression results for unique sources per subscriber (sqrt transformed)

```
Fixed-effects (within) regression               Number of obs      =       664
Group variable (i): opcode_new                  Number of groups   =       175

R-sq:  within  = 0.0765                         Obs per group: min =         1
       between = 0.1200                                        avg =       3.8
       overall = 0.1029                                        max =         4

                                                F(8,481)           =      4.98
corr(u_i, Xb)  = -0.9999                        Prob > F           =    0.0000

------------------------------------------------------------------------------
  src_per_sq |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
   totsub_ln |  -.0546212    .028605    -1.91   0.057    -.1108273    .0015849
   srv_cable |  (dropped)
   cblXsubln |  -.0044461   .0460883    -0.10   0.923    -.0950054    .0861131
     lap_mem |  (dropped)
 piracy_rate |   .1078349   .0472453     2.28   0.023     .0150023    .2006675
     educ_ix |   32.29357   6.626307     4.87   0.000     19.27349    45.31366
     lapXedu |  -34.13204   8.034932    -4.25   0.000    -49.91995   -18.34414
     lapXpir |  -.0877637   .0738087    -1.19   0.235    -.2327911    .0572637
     eduXpir |  -.1185407   .0495059    -2.39   0.017    -.2158153    -.021266
 lapXeduXpir |   .0926636   .0781651     1.19   0.236    -.0609236    .2462509
       _cons |  -10.84459   3.551544    -3.05   0.002    -17.82305   -3.866135
-------------+----------------------------------------------------------------
     sigma_u |  15.753836
     sigma_e |  .09650475
         rho |  .99996248   (fraction of variance due to u_i)
------------------------------------------------------------------------------
F test that all u_i=0:     F(174, 481) =    10.25                 Prob > F = 0.0000
```

approach for finding a statistical model that yields statistically significant parameter estimates. Table 4 presents the findings from a version that is close to the one found to fit the pooled data.[7] Several variables (number of subscribers, interaction terms) show the same sign as in the other procedures. The coefficient for the education proxy in the panel approach shows a positive sign, which is incompatible with the more plausible findings in the other two approaches. In the case of total subscribers, the significance level is below the 5 percent level. Statistical significance may not be a central concern, though, because the ISPs in our dataset represent the lion's share of the ISP market in the 40 nations. Therefore, our data is nearly a complete enumeration of the markets rather than a sample. In this case, the parameter estimates reflect the data structures in the empirical universe under investigation. In as far as ISPs are concerned, it is not necessary to make inferences from a subset to the whole phenomenon. Consequently, significance levels lose in importance when interpreting the findings.

[7] The model explains only about 10 percent of the variance in the dependent variable overall, 12 percent of the variance between ISPs, and 7 percent of the variance within each of the 175 ISPs for which all data were available. Given the wide diversity of the ISPs, the short panel of only four years of observations, and the highly dynamic phenomenon of spam, this is not surprising. Other model specifications explain a higher share of the variance but often do not yield statistically significant coefficients. If we were to interpret the data as an enumeration, this would be acceptable. In these model runs, we can explain up to 49 percent of the overall variance of the total number of spam messages, with the size of the user base typically the most important explanatory variable.

Table 5: Summary empirical results for unique sources per subscriber*

	bivariate	pooled	panel
total_sub_log	negative	negative	negative
srv_cable	cable lower	negative	---
lap-mem	negative	positive	---
cyberconv_mem	negative	negative	---
piracy_rate	positive	negative	positive
edu-ix	negative	negative	positive
lapXedu	---	negative	negative
lapXpir	---	negative	negative
lapXpirXedu	---	positive	positive

* Shaded cells mark consistent results across different approaches.

Table 5 summarizes and compares the findings from the different approaches. Explanatory variables which show the same direction of influence on the dependent variable in all three approaches can be considered more robust than independent variables for which the findings differ (shaded in Table 5). In several cases, a variable could not be used in all three specifications, for example, because it did not vary within on country during the four years under consideration. Moreover, interaction terms do not make sense in bivariate statistics. In these cases, consistency means that the variable show the same directional effect in the remaining two approaches. Such defined consistency is observable for the size of ISPs, for whether an ISP uses a cable rather than a DSL platform, membership in the cybercrime convention, and the various interaction terms between LAP membership and education as well as LAP membership and piracy. In the case of the piracy rate and the education proxy results are mixed, with one method's results deviating from the others. We interpret this as a less robust finding.

Conclusions

This paper set out to address a number of questions. First, our findings support the view that ISPs are indeed critical control points for botnet mitigation. In addition, a more specific pattern was uncovered. While the class of ISPs includes anywhere between 4,000 and 100,000 actors, we found that the distribution of infected machines is highly asymmetrical. Just 50 ISPs consistently accounted for over half of all infected sources. Such a skewed distributions is a familiar pattern for many Internet-related phenomemon. However, its presence in this situation is less likely than it may appear, as ISPs for consumers and SMEs are oriented towards national markets, not global ones.

From a policy perspective, this is a relevant finding. Even if ISPs were to be a more effective control point compared to the hundreds of millions of end user machines, it would be extremely difficult to bring about collective action among many thousands of actors located in over a hundred countries. Our data suggests the task may have more manageable proporties. Not only is the number of actors needed to create an impact on botnets smaller than expected, the most critical actors are also the easiest to target with governmental interventions or some form of public-private sector cooperation, as they are

larger, well-established corporations, rather than large numbers of small ISPs that are often shortlived and difficult to survey, let alone reach with collaborative or regulatory efforts.

Stimulating ISPs to increase botnet mitigation presumes that ISPs have the discretion to step up such efforts. This is not self-evident. It is well-known that in retail ISP markets, competition is first and foremost driven by price. In many countries, price competition is fierce. Moreover, even if consumers cared about security, there are no adequate market signals that could reliably guide them towards more security-conscious ISPs. Most industry insiders lack such signals as well, except for the unreliable anecdotal evidence and speculative claims that are bandied around the security community about the performance of this or that ISP.

If the behavior of ISPs is mostly driven by institutional incentives, outside the control of the individual ISP, then we would expect similar levels of performance in terms of botnet mitigation. Attempts to get ISPs to increase their efforts would first have to change that incentive structure. To get a sense of the discretionary power of ISPs to do botnet mitigation, we explored the extent in which they performed different relative to each other, in terms of the number of infected machines in their networks. We found that perfomance levels are highly dispersed. For ISPs of similar size, we found that the differences typically span two orders of magnitude – i.e., a hundred-fold difference. Even within the same country, we see differences of more than one order of magnitude for ISPs of similar size. In other words, external conditions do not dictate the ISPs' internal incentives and, hence, their efforts. Operating under comparable conditions allows for remarkable differences in performance.

We developed a theoretical framework to explain the differences among ISPs and then empirically tested some of these explanations. We found that characteristics of the user base matter. Higher rates of using pirated software are associated with higher botnet activity. Higher average connection speeds are not. The level of education, as a proxy for technical competence, is associated with lower levels of botnet activity. We also found limited evidence to support the idea that governmental efforts to improve cybersecurity are related to lower levels of botnet activity – confirming earlier research Png and Wang (2007) and Wang and Kim (2009), though unlike the latter, we found no impact of the Convention on Cybercrime, once we took other factors into account. However, given the substantial variability among ISPs subject to one specific set of institutional incentives, such public policy measures, while possibly necessary conditions to enhance security, are, taken by themselves, not sufficient.

Regarding the ISPs themselves, we found that average revenue per customer did not make a difference. So price may not be related to security performance. Market share of an ISP in its home country was not associated with worse performance either. We also tested the claim that large ISPs perform worse than smaller ones, because they are less subject to peer pressure (Moore et al. 2009). Our data suggests this is incorrect. In fact, we found support for the idea that large ISPs actually perform better than average, measured in number of sources and spam volume per subscriber). The reason that

industry insiders often claim the opposite may simply be an effect of repeatedly seeing the same names on the lists of worst offenders – in other words, of seeing the 50 or so ISPs that we also found as critical to the overal problem. However, those lists may fit nicely with anecdotal evidence, but they fail to take into account critical and obvious factors, such as the size of the customer base. One speculative reason why large ISPs actually do slightly better may be that their size forces them do introduce automation in incident response and abuse management. A similar mechanism may explain why we found that cable providers did slightly better than DSL providers. The management of cable networks often include automated systems and these technologies perhaps make it less costly to deal with infected machines. Given the ongoing advances in technology, including botnet mitigation solutions, the difference between cable and DSL may disappear in the immediate future.

In sum: our study provides evidence that ISPs are critical control points and that even under current market conditions increased efforts to mitigate botnets appear possible. Current efforts to bring about collective action – through industry self-regulation, co-regulation, or government intervention – might initially achieve progress by focusing on the set of ISPs that together have the lion's share of the market. Further work is needed to explore ways in which to strengthen the ISPs incentives to improve botnet mitigation.

Appendix 1: List of the countries and count of ISPs included the final dataset

Code	Country Name	OECD status	Number of ISPs
AT	Austria	Member	3
AU	Australia	Member	6
BE	Belgium	Member	4
BR	Brazil	Enhanced engagement	8
CA	Canada	Member	9
CH	Switzerland	Member	3
CL	Chile	Candidate	5
CN	China	Enhanced engagement	5
CZ	Czech Republic	Member	4
DE	Germany	Member	13
DK	Denmark	Member	3
EE	Estonia	Candidate	2
ES	Spain	Member	6
FI	Finland	Member	4
FR	France	Member	5
GB	United Kingdom	Member	8
GR	Greece	Member	3
HU	Hungary	Member	6
ID	Indonesia	Enhanced engagement	2
IE	Ireland	Member	7
IL	Israel	Candidate	3
IN	India	Enhanced engagement	6
IS	Iceland	Member	2
IT	Italy	Member	4
JP	Japan	Member	6
KR	South Korea	Member	4
LU	Luxembourg	Member	1
MX	Mexico	Member	5
NL	Netherlands	Member	6
NO	Norway	Member	5
NZ	New Zealand	Member	4
PL	Poland	Member	5
PT	Portugal	Member	4
RU	Russia	Candidate	10
SE	Sweden	Member	4
SI	Slovenia	Candidate	5
SK	Slovakia	Member	2
TR	Turkey	Member	1
US	United States	Member	15
ZA	South Africa	Enhanced engagement	2
TOTAL	**40**		**200**

Appendix 2: Data and data sources

Category	Variable	Description	Source
Dependent variables	*unq_srcs*	Number of unique IP sources emitting spam from an ISP during a specific time period.	Processed spam trap data
	spam_msgs	Total number of spam messages (spam volume) emitted from an ISP during a specific time period.	
	unq_srcs_sub	Unique sources per subscriber. Similar to unq_srcs, but corrected for size of the ISP	
	spam_msgs_sub	Spam messages per subscriber. Similar to spam_msgs, but corrected for size of the ISP	
Independent variables	*total_subs*	Total number of subscribers of an ISP (retail, business, DSL, cable, etc)	TeleGeography GlobalComms
	srv_type	The type of service / access provided by the ISP: DSL, cable, or both. A variant of this variable is *srv_cable* (1 if ISP provides cable access).	
	rev_per_sub	Revenue of the ISP (wireline section) divided by its subscriber count.	
	int_bpp	International Internet bandwidth, per person, in the country the ISP operates in (measured in bits per person).	World Development Index
	bb_subs	Number of broadband Internet subscribers in the country the ISP operates in. (this variable is used indirectly, in calculating *market_share*)	
	lap_mem	Is the country in which the ISP is located, a member of the London Action Plan?	Own construction
	cyberconv_mem	Has the country in which the ISP is located, signed the convention on cybercrime?	
	piracy_rate	Percentage of software that is pirated in the country the ISP operates in.	Business Software Alliance
	educ_ix	Education index: an index indicating the overall education level of people in the country that the ISP operates in.	UN Human Development Reports
	market_share	Local market share of the ISP (*total_sub* divided by *bb_subs*)	TeleGeography GlobalComms
Mappings	*ASN-to-AS-name*	Mappings of Autonomous System numbers to names	WHOIS
	AS-name to ISP	Mappings of ASNs to the ISPs (i.e., which ISP owns which ASN)	Own construction
	ASN to country	Mappings of IP addresses to countries (IP location)	MaxMind GeoIP

Appendix 3: Descriptive statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
unq_srcs	741	185324.9	504712.2	20	5904500
src_persub	741	.191359	.2107284	.0001	1.1329
total_sub	741	1374231	3369101	3000	4.43e+07
market_share	709	.1820523	.1988902	.0005	1.2358
rev_persub	194	4305.685	5135.762	182.1285	42768.34
srv_cable	665	.3233083	.4680914	0	1
lap_mem	741	.562753	.4963815	0	1
cyber_mem	741	.7098516	.4541373	0	1
piracy_rate	740	40.35135	17.31936	20	87
educ_ix	741	.9440931	.0671315	.632	.993
int_bpp	386	14345.77	16918.11	190.8559	92832.46
spam_msgs	741	4.53e+07	1.10e+08	7679	1.54e+09
spam_persub	741	55.57274	79.03818	.1697	830.8522

Appendix 4: Pair wise correlations between the independent variables

	total_~b	market~e	rev_pe~b	srv_ca~e	lap_mem	cyber_~m	piracy~e
total_sub	1.0000						
market_share	0.2530	1.0000					
rev_persub	-0.1149	0.1907	1.0000				
srv_cable	-0.0958	-0.2110	-0.2721	1.0000			
lap_mem	0.1519	-0.1900	-0.1078	0.0922	1.0000		
cyber_mem	-0.0729	-0.0887	0.0188	0.0738	0.1498	1.0000	
piracy_rate	0.0849	0.1284	0.1344	-0.0942	-0.3513	-0.6066	1.0000
educ_ix	-0.0835	-0.0793	-0.2180	0.1317	0.3184	0.4939	-0.6068
int_bpp	-0.0527	-0.0759	-0.1881	0.0046	0.2420	0.4757	-0.5176

	educ_ix	int_bpp
educ_ix	1.0000	
int_bpp	0.3476	1.0000

References

Anderson, R., R. Böhme, R. Clayton and T. Moore (2008). *Security Economics and the Internal Market*. ENISA (European Network and Information Security Agency). Available online at http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf.

Asghari, H. (2010). *Botnet mitigation and the role of ISPs: A quantitative study into the role and incentives of Internet Service Providers in combating botnet propagation and activity*. Faculty of Technology, Policy and Management. Delft University of Technology. Available online.

Bauer, J. M. and M. Van Eeten (2009). Cybersecurity: Stakeholder incentives, externalities, and policy options. *Telecommunications Policy* 33(10-11): 706-719.

Bauer, J. M., M. J. G. Van Eeten and T. Chattopadhyay (2008). *ITU Study on the Financial Aspects of Network Security: Malware and Spam*. ITU (International Telecommunication Union). Available online at http://www.itu.int/ITU-D/cyb/cybersecurity/docs/itu-study-financial-aspects-of-malware-and-spam.pdf.

BBC News (2007). *Google searches web's dark side*. BBC News website. Available online at http://news.bbc.co.uk/2/hi/technology/6645895.stm.

House of Lords (2007). *Science and Technology Committee, 5th Report of Session 2006–07, Personal Internet Security, Volume I: Report*. Authority of the House of Lords. Available online at http://www.publications.parliament.uk/pa/ld/ldsctech.htm.

Ironport (2006). *Spammers Continue Innovation: IronPort Study Shows Image-based Spam, Hit & Run, and Increased Volumes Latest Threat to Your Inbox*. Available online at http://www.ironport.com/company/ironport_pr_2006-06-28.html.

Jakobsson, M. and R. Zulfikar, Eds. (2008). *Crimeware: Understanding New Attacks and Defenses*, Addison-Wesley Professional.

Messagelabs (2009). *MessageLabs Intelligence: Q3/September 2009*. Available online at http://www.messagelabs.com/resources/mlireports.

Moore, D., C. Shannon and J. Brown (2002). *Code-Red: a case study on the spread and victims of an Internet worm*. Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment. Available online at http://portal.acm.org/citation.cfm?id=637244.

Moore, T., R. Clayton and R. Anderson (2009). The Economics of Online Crime. *Journal of Economic Perspectives* 23(3): 3-20.

OECD (2009). *Computer Viruses and Other Malicious Software*. Paris, Organisation for Economic Co-operation and Development.

Png, I. and C. Y. Wang (2007). *The Deterrent Effect of Enforcement Against Computer Hackers: Cross-Country Evidence. Paper presented at the 2007 Workshop on the Economics of Information Security.* Available online at http://weis2007.econinfosec.org/papers/77.pdf.

Spindler, G. (2007). *Verantwortlichkeiten von IT-Herstellern, Nutzern und Intermediären: Studie im Auftrag des BSI durchgeführt von Prof. Dr. Gerald Spindler, Universität Göttingen*. Bundesamt für Sicherheit in der Informationstechnik. Available online at https://www.bsi.bund.de/cae/servlet/contentblob/486890/publicationFile/30962/Gutachten_pdf.pdf.

Stone-Gross, B., M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel and G. Vigna (2009). *Your Botnet is My Botnet: Analysis of a Botnet Takeover*. 16th ACM Conference on Computer and Communications Security, November 9–13, 2009, Chicago, Illinois. Available online at http://www.cs.ucsb.edu/~seclab/projects/torpig/torpig.pdf.

US GAO (2007). *Cybercrime: Public and Private Entities Face Challenges in Addressing Cyber Threats*. United States Goverment Accountability Office. Available online at http://www.gao.gov/new.items/d07705.pdf.

Van Eeten, M. and J. M. Bauer (2008). *Economics of Malware: Security Decisions, Incentives and Externalities, OECD STI Working Paper 2008/1*. OECD. Available online at http://www.oecd.org/dataoecd/53/17/40722462.pdf.

Wang, Q.-H. and S.-H. Kim (2009). *Cyber Attacks: Cross-Country Interdependence and Enforcement. Paper presented at the Eighth Workshop on the Economics of Information Security (WEIS 2009).* Available online at http://weis09.infosecon.net/files/153/paper153.pdf.

Zhuang, L., J. Dunagan, D. R. Simon, H. J. Wang, I. Osipkov, G. Hulten and J. D. Tygar (2008). *Characterizing Botnets from Email Spam Records*. LEET '08. First Usenix Workshop on Large-Scale Exploits and Emergent Threats, San Francisco. Available online at http://www.usenix.org/event/leet08/tech/full_papers/zhuang/zhuang.pdf.