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

Master Thesis
by Hadi Asghari (1393707),
MSc. Management of Technology

Delft University of Technology
Faculty of Technology, Policy and Management
Section Policy, Organization, Law and Gaming

January 2010

Thesis Committee:

Chairman and first supervisor: Prof. Michel van Eeten (Professor of Public Administration, Section POLG)
Second supervisor: Dr. Roland Ortt (Associate Professor of Technology Management, Section TSE)
Third supervisor: Andreas Schmidt (PhD Candidate, Section ICT)
External expert: Dave Rand (Chief Technologist of Internet Content Security, Trend Micro Inc.)

Dedicated to those who empower the citizen media of the 21st century

ACKNOWLEDGEMENTS

Personal acknowledgements

Writing this thesis would not have been possible without the support of several people whom I need to wholeheartedly thank.

If I were to create a list, it would without doubt be topped by my parents. They have given me constant encouragement, support and inspiration in obtaining higher education, and also made my studies in TU Delft possible in the first place. There wish is for their children to seek knowledge that benefits society, and I hope that my work will do so.

During my current studies, I have been accompanied by my beloved wife, Golrokh. I know that the choice of being far from the place she calls home has not always been easy for her - particularly during these special times for Iran and apart from friends. Her support has been unique in days which I put all my energy in completing my research.

I have also been lucky enough to have worked under the supervision of Michel and Roland. Both of them are excellent researchers who have inspired me by their high standard of work, and more importantly, provided me constant guidance in my academic work, and helped in many other ways.

Of course, this list is not exhaustive in anyway. Many other friends, in both Iran and the Netherlands, assisted me in my journey - towards each of whom I am always grateful.

Hadi Asghari, January 2010, Delft.

Academic acknowledgements

This study is part of a larger research project being executed for the Organisation for Economic Co-operation and Development on the same subject. As such, this thesis borrows many concepts and ideas from this larger project, and it is often hard to point them out individually. The OECD project is headed by Prof. Michel van Eeten (my first supervisor), and other members of the team include Prof. Johannes Bauer (professor at Michigan State University), Shirin Tabatabaie (PhD candidate at TU Delft) and Woohyun Shim (PhD candidate at MSU). Certain contributions to this work that can be specifically named are as follows: Prof. Van Eeten suggested the original topic; Ms. Tabatabaie played an important role in preparation of the datasets, specifically in gathering the country level variables, a painstaking task; and Prof. Bauer provided valuable advice in the execution of the data analysis section. Additionally, Menno Nederveen, a student assistant at the section, provided help in the ASN to operator mappings.

Dave Rand of Trend Micro Inc also needs to be separately and sincerely thanked. Dave generously shared with our group the exclusive spam data that has been used as the basis of the analysis.

EXECUTIVE SUMMARY

This thesis is about *botnets*. Botnets are networks of compromised computers, that unknown to their owners, run a malicious piece of software (called a *bot*). This code puts the computer under the control of a remote attacker, who then uses these bots to accomplish a variety of illegal tasks - from sending spam emails and disruption of the network, to identity and financial theft. Botnets are a serious threat. Exact numbers are hard to come by, but what is certain is that millions of computers worldwide have been compromised (estimates are 1% to 25% of online users). Criminals make "real" money off these bots, and of course, inflict very real damage – estimated in billions of euro; what's more, the overall trends appear to be worsening.

Thus this thesis is more correctly about *mitigating* botnets. Combating botnets is not an easy task, due to the combination of factors that make them possible in the first place. Foremost, it is very lucrative for criminals to 'recruit' bots. They do this by exploiting software vulnerabilities, and by tricking users to run them. Unfortunately, there seems to be no shortage of software vulnerabilities (- this is not a coincidence!). Additionally, given the dynamic nature and complexity of the attacks, a portion of users will be unable to protect themselves. An interesting idea that has been floating around for some time is to have the *Internet Service Providers* act as 'gatekeepers', and help protect users residing on their networks. The rationale is that ISPs can notice unusual traffic patterns being exhibited by their subscribers. For example, attempts to access known 'Command and Control' domains is a tell-tale sign of a bot infection. ISPs could then notify the subscriber, and even disconnect her until the infection is removed. A beautiful and practical solution; but the ISPs have little incentive to implement it, particularly on a large scale.

It should be obvious at this point that we are looking at a socio-technical problem. Applying economic theory, we can describe the problem in terms of the *misaligned incentives* of a network of actors that creates *negative externalities* for all. Based on prior studies, we know that ISPs face two sets of costs in regards to infected subscribers. They have incentives to take *some* action against bots (to avoid the costs of blacklisting, losing reciprocity, etc); but not *too much* action (to avoid costs of customer calls, legal liability, etc). In balancing these costs, most ISPs decide to take action against only the most aggressive bots on their networks. This implies that ISPs will differ on the degree of vigilance they adopt against bots - based on the strengths of their various incentives. If we could somehow measure the strengths of these incentives, we might be able to design mechanisms to mitigate botnets, via the ISPs.

Thus, our research question was born, as follows: "Are ISPs crucial intermediaries in botnet mitigation efforts? Do they significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? And what implications does this have for policy?" To answer this question, we set out to quantify the degree in which ISPs mitigate botnets, and connect this with economic variables regarding the ISP itself (such as number of subscribers), and regarding the country in which it operates (such as a regulation index). The mitigation efforts of ISPs were measured by using *outbound spam as a proxy for botnet activity*. This proxy works because in the past few years spam has been sent mostly via infected machines.

A large and unique dataset containing the list of IP addresses of spam sending machines, each day, for the years 2005 to 2008 was used for this purpose. (The data was logged by a 'spam-trap', and checked to be a *representative* sample). Using this sample, the number of infected machines within each ISP worldwide was determined. By dividing this count by the number of subscribers within each of them, we arrive at our dependent variable - a measure of relative botnet activity within each ISP. Building this variable, although methodologically straight forward, necessitated the setup of a computing infrastructure capable of processing *over one billion* log records.

(The infrastructure consists of a 'high performance cluster' running self-developed Python scripts, and has the potential of being used in similar research.)

The independent variables come from a variety of sources, one of which is a commercial database of market information on retail ISPs worldwide. To aid ourselves in the selection of the independent variables, a two stage approach was used. First, a *conceptual framework* was developed by synthesising the existing literature, on what influences botnet activity at the ISP level. (The model is basically as follows: criminal activity and user behaviour are the main causes of botnet activity; this casual relation is intensified by technological enablers, and mitigated by security measures of ISPs. The ISPs choose their measures based on their mix of incentives and cost perception). Testing all the relations in the model was not possible due to unavailability of data, so in the second stage, a *measurement model* was developed – an adaptation of the theoretical model based on empirical data obtainable.

A set of nine empirical hypotheses were then formulated (the literature being very scarce), and statistically tested. The data analysis was performed in two steps: an individual testing of hypothesis, supplemented by a multivariate regression analysis. The final dataset contains 741 cases, representing four years of observations for the security performance of 200 ISPs. These are the major retail ISPs operating in the 40 countries of the *enlarged-OECD*. The major findings of the hypothesis tests are as follows:

- The results provide *compelling evidence that ISPs are in fact the focal point in botnet mitigation*: approximately 200 retail ISPs account for 80% of the bot infections in these countries; in other words, the bulk of the problem is concentrated within a small number of economic players.
- Evidence is also found that retail brand ISPs differ significantly in regards to the level of relative botnet activity occurring on their networks. Variability among ISPs of similar size suggests that the *choice of security practices ISPs adopt makes a big difference.*
- Another major finding is that *the count of subscribers is negatively associated with botnet activity levels*. This contradicts the commonly held belief that larger ISPs perform worse in terms of security.
- Interestingly enough, *cable providers have a better security performance then DSL providers* – on average 10% lower bot infections. (Some reasons for this could be the existing use of traffic monitoring systems in such networks, due to the shared bandwidth infrastructure; alternatively, it could be due to the possibility to enforce stricter network policies, as they typically have a large residential subscriber base.)
- *Targeted regulation such as those stimulated by the LAP appear to be effective* - ISPs operating in countries that have signed the 'London Action Plan' have on average 13% lower bot infections. Conversely, the broader 'Convention on Cybercrime' appears to be ineffective.

Other results include finding *piracy rate* to be positively associated with botnet activity levels; and *ARPU* and *market share* to seemingly have no influence on botnet activity.

The general word of caution regarding bivariate relations (that found associations might disappear after controlling for other variables) led us to also perform multivariate regression analysis. This revealed interesting results too:

- Approximately *40% of the sample variance* regarding the relative degree in which ISPs mitigate botnets (i.e., number of infected sources corrected for size), can be explained using the variables *subscriber count, cable access, LAP membership, privacy rate, education index, and year*. (This percentage is high, considering the limited number of variable used in explaining a complex phenomenon)

- The interaction terms among the country level variables (LAP membership, piracy rate, and education index) in the regression model indicate that these variables *change and find meaning in configurations* – often the case with demographic and institutional variables.
- An *interaction term also exists between subscriber count and cable access*. The existence of this term, coupled with the direction of the beta, strengthens speculation that the increased security performance in these organizations has a common cause – i.e., the use of automated abuse monitoring and handling systems in large ISPs (due to scalability) and in cable providers (as already explained).

When searching for the policy implications of these findings, and by taking a look back at the research question, we can state that: yes, ISPs are crucial intermediaries in botnet mitigation efforts (and *can* act as the gatekeepers); and yes, they differ significantly in the degree in which they mitigate botnets (which provides policy makers room for starting negotiations with them).

Among the factors investigated, the most promising are that targeted regulation seems to be effective; and the fact that larger retail ISPs and cable providers have lower botnet activity levels (when corrected for size). This prompts us to search for the underlying case - if the speculation that it is due to the use of automated abuse monitoring and handling systems turns out to be true, then the adoption of such systems should be encouraged.

However, critical question need to be answered before following these recommendations. One is whether in the broader context the situation of cyber-security will improve, should the ISPs be made partially responsible for mitigating botnets - given that the root causes for cyber-insecurity still remain. Another is whether society will be better after such an initiative is implemented – given the possibility of opportunistic behaviour (by ISPs, as well as others), and given the trajectories that the initiative might create. Plausible arguments have been expressed indicating that the overall situation might actually end up being worse. Balancing these arguments is in the end however a political decision, and must be done by the policy makers.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1 – RESEARCH PROPOSAL

1.1 RESEARCH CONTEXT [1]

Over the past decade, the Internet has transformed into a critical infrastructure, with millions of citizens and businesses in developed countries depending on it on a daily basis. During this same period, the threats facing Internet security have also increased. The same characteristics that made the Internet economy grow so successfully – such as the ability for every node to run arbitrary code – are being abused by criminals, to commit illegal activities on a vast scale. The global costs of such activities are in the magnitudes of billions of euro (Bauer, Van Eeten, and Chattopadhyay 2008). It is evident that measures must be taken to protect the Internet from such threats.

This research aims to contribute to the scientific debate in the fields of Internet Security and Economics of Information Security. Before going any further, it needs to be stated that the concept for this research has been born from the previous works of the author's first supervisor, Prof. Van Eeten. It forms part of a bigger research project on the same subject being conducted for the OECD, by Prof. Van Eeten, Prof. Bauer[2], other colleagues, and this author.

But what makes Internet security an interesting topic for social scientists? A major reason is that technical solutions alone fail to produce satisfactory results. Take the case of *spam emails*. Over the past few years, security experts have come up with many innovative technical solutions to do away with spam once and for all[3], but none of them have managed to do so. Instead, spam rates have steadily increased to an astonishing 120 billion spam emails daily in 2008 (IronPort 2008b)[4]. The spammers have simply evolved their techniques, creating a technological arms race. The same story holds for other online threats, such as *malware*, *phishing*, and *DDoS*. They are all growing in sophistication and strength.

Among all the threats, *botnets* are without doubt the hardest to tackle (as we shall see in later chapters of this thesis). A *bot* is a malicious piece of software that runs on a compromised computer system, without its owner's knowledge, and in cooperation with other bots, accomplishes some illegal tasks. Figure 1 shows a botnet: the computers marked with 'Z' are under the control of the bot-master, unknown to their owners. Botnets are very profitable to run for criminals, as they are the platform to launch and host many network attacks. It is estimated that millions of computers have been "recruited" into botnets (BBC 2009d). No clear-cut solution exists for mitigating botnets, as they are the result of many socio-technical shortcomings.

'*Misaligned incentives'* is one of the perspectives used to explain the status quo of Internet security. Basically put, rational decisions that legitimate actors make regarding their own security , which is influenced by their individual

[1] Please note that this chapter is a modification of the originally submitted research proposal. Some parts, including the research questions, were updated as the project progressed along.

[2] Prof. Johannes Bauer is a professor of Telecommunication, Information Studies, and Media, and the Director of Special Programs at Quello Center for Telecommunication Management & Law, in Michigan State University.

[3] For instance, spam filters, CAPTCHAs, authentication systems, IP reputation, etc.

[4] A word of caution must be raised that security firms have incentives to over-report such figures. Nonetheless, they remain one of the most important sources of statistics on cyber-crime and online threats. A wide range of such statistics will be examined in later chapters, and it's important to know that even the most conservative reports agree on the worsening of the situation.

incentives, do not add up to a very secure Internet; tradeoffs are not sufficiently aligned with the social-optimum. In this perspective, understanding and fine-tuning the incentive structure can be a powerful and sustainable approach for combating botnets.

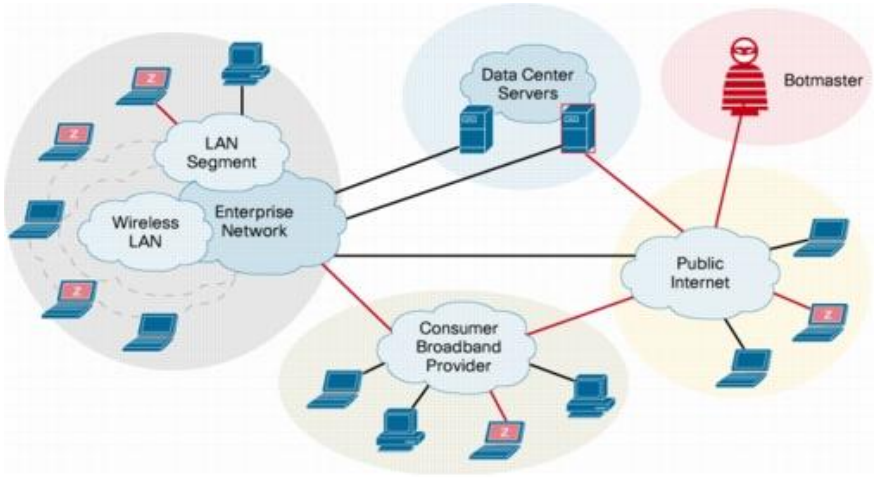

Figure 1 – A typical botnet – The computers marked with Z are bots (Cisco 2007)

A hotly debated topic these days is whether Internet Service Providers should do more to fight botnets. Proponents of this view believe that focusing on ISPs is the most practical solution to botnets, as ISPs are in a position to block malicious traffic (originating from bots) from leaving their networks, thus keeping the Internet "clean". They also state that by monitoring their networks, ISPs can identify infected machines and help customers in remediation. Figure 2 illustrates these measures. Some advocates go so far as to recommend holding ISPs liable for a failure "to respond promptly to requests for the removal of compromised machines" (Anderson et al. 2008b). Critics of this view also raise several points, most importantly that ISPs cannot economically cope with today's scale of bot infections, should they be held liable(Economist 2009). Van Eeten and Bauer (2008) believe that ISPs already have incentives to pursue infected machines, but only to a certain degree. The behaviour of ISPs, like all rational agents, is determined by balancing the costs against the benefits of adopting particular security measures.

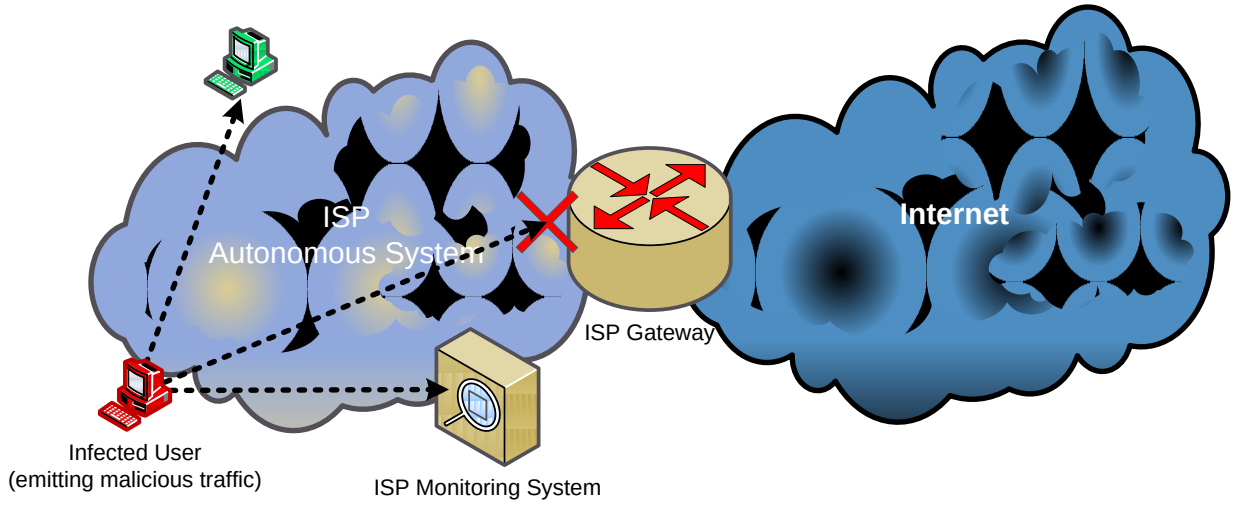

Figure 2 - Examples of how ISPs can detect and filter bot traffic

We will take a more detailed look at this debate in the next chapters, but suffice it to say that this 'certain degree' is rather broad: on one extreme, we have ISPs who can remove an infected machine in under one hour, on the other end, we have the *"rogue"* ISPs, who virtually take no action; with the majority lying somewhere in between these two extremes (Anderson et al. 2008b; Van Eeten and Bauer 2008). In this research we aim to empirically and quantitatively examine the mentioned difference between ISPs. If we can identify the magnitude of the difference, and understand the underlying incentives and factors, we will have taken an important step in the debate regarding botnet mitigation.

A unique opportunity for this research exists in the TPM faculty: access to an exclusive and massive dataset of malicious traffic, emitted from ISPs worldwide over a period of four years. This dataset can be used to measure ISP security effectiveness, and combined with other variables, used as a basis for our analysis.

1.2 Research objective and questions

Botnets are one of the most serious threats the Internet faces today, and like other cyber security challenges, they cannot be eradicated by technical solutions alone. Rather, the underlying economic incentives of all actors involved must also be addressed. Among these actors, the role of Internet service providers in combating bots has been the topic of much controversy lately. Our research objective is to investigate this controversy. Since ISPs have great variation in how they react to infected machines on their networks, a good starting point will be trying to understand these existing differences between ISPs, and what they are caused by. This leads up to the following problem statement:

> ***Problem statement:*** *Are Internet Service Providers crucial intermediaries in botnet mitigation efforts? Do they significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? And what implications does this have for policy?*

To answer this question, we would need to answer the following sub questions:

> ***SubQ1:*** *What are botnets and why is it important to mitigate them?*
>
> ***SubQ2:*** *Who are the main actors that can mitigate botnets? Are the ISPs central?*
>
> ***SubQ3:*** *Do ISPs significantly differ in the degree to which they mitigate botnets?*
>
> ***SubQ4:*** *To what extent can we explain the varying degree in which ISPs mitigate botnet activity? Can we identify internal or external factors that can explain this variance?*[1]
>
> ***SubQ5:*** *What are the implications of the above findings in terms of policy options?*

1.3 Scientific background

There are multiple scientific frameworks underlying our research. Foremost, our research builds upon a previous paper by the supervisor of the applicant, titled "Economics of Malware". That research used a *qualitative* empirical approach to examine the decisions, incentives, and externalities caused by various market players in the Internet economy. Our research takes a *quantitative* empirical approach and examines only one intermediary actor.

[1] This question has both quantitative and qualitative aspects

But taking another step back, our research is part of a new discipline called *"economics of information security"*, which came to the intellectual attention around 2000 (Wikipedia 2009c). One ground breaking paper was the work of Ross Anderson (2001), entitled *"Why Information Security is Hard"*. The paper described the history of the ATM machines in Europe and United States. Although the technology used was approximately equal in both regions, different liability laws (surrounding fraud) caused a totally different security landscape to emerge.

Basically, in the US, in cases of fraud, the bank has to prove that the user was at fault, or otherwise take the responsibility. In Europe, the reverse is true – users have to prove that the bank was at fault – a virtually impossible endeavour. After 20 years, the results are astonishingly different: levels of fraud in the US are much lower than Europe. The argument is that banks (which have the power to actually influence ATM security), when faced with the correct incentives (e.g., fraud liability), take concrete steps in increasing security.[1]

Other recent and relevant influential work has been Shapiro & Varian's *"Information Rules: A Strategic Guide to the Network Economy" (2000)*. These authors argued that although the internet is a totally new medium, the economical rules governing it have been around for quite a long time: network effects, information asymmetry, lemon markets, etc. For example, software is not as secure as it should be, because the first mover advantage outweighs the costs of being seen as unsecure and providing patches in a later stage. In this sense, the software vendors are externalizing their costs onto others.

1.4 Research approach

The theoretical framework that we shall use is presented in Figure 3. (This framework is developed and described in the literature review chapter).[2] Based on this framework, empirical hypotheses will be stated, in the form of relations which we expect to hold regarding the relations between various incentives, factors, and botnet activity.

This data for measuring botnet activity will come from a *spam trap* operated by Dave Rand, an international security expert. We shall use spam as a proxy for botnet activity and propagation. We believe this is a valid proxy as more than 90% of spam today is sent via botnets (MessageLabs 2009), and an IP address sending spam most likely points to an infected machine. Each record in this dataset consists of the *source IP address* and *source ASN (autonomous system number)* of a spamming machine, in addition to the date the spam was received. Using these variables, we can deduct the geographic location and the originating ISP.

This spam trap has logged more than 60 billion spam messages, originating from over 21,000 networks since 2005. To process such a huge dataset, and aggregate it to levels suitable for statistical software, the *Python* scripting language will be used. Other variables for the model need to be gathered from other sources, such as ISP market data, and country level statistics.

A final point worth mentioning is the necessity to *triangulate* our dataset before starting analysis. This is to make sure that the dataset provided to us is a representative sample of worldwide spam activity. One of the ways to do this is to compare our aggregated country statistics with figures reported publicly by security firms.

[1] Examples include training personnel, installing cameras, detecting fraud patterns, etc

[2] As stated on the first page, note that the framework in the initial research proposal differed from the one presented here.

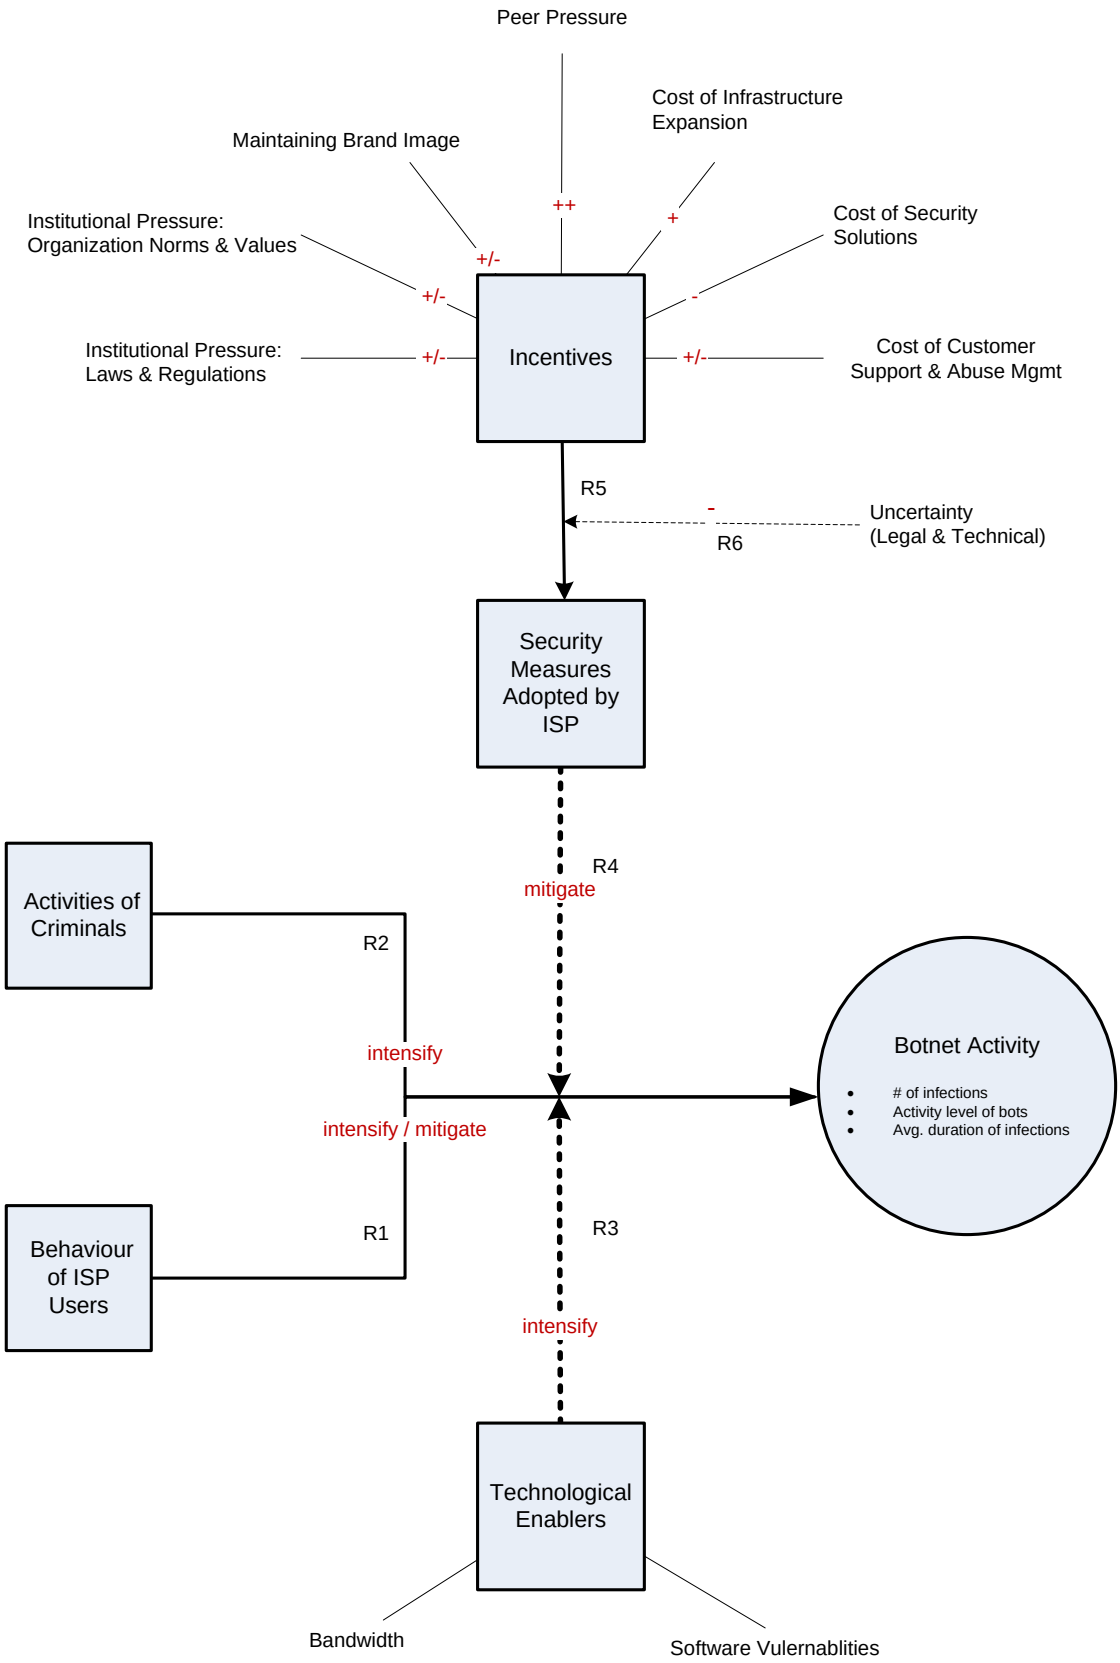

Figure 3 - Theoretical framework

1.5 Scientific and management relevance

This research fits in with current on-going discussions in the "security economics" literature (as explained in the scientific background section). The added value of our work will be in the empirical contribution and quantitative analysis, which is at this time missing. In a recent report, Anderson and his colleagues (2008b) recommended that the European Network and Information Security Agency *"collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs"*, adding that *"once decent statistics are available, the next question is what further incentives might help"*. This is precisely the research gap we are hoping to fill.

From a management perspective, a great sense of urgency exists to tackle the botnet problem, and as we argued, technical means alone are not capable of solving this problem. Policy makers have great interest in identifying the underlying incentives and influencing factors attributing to botnet proliferation, in order to come up with economical and legal measures that can reduce them. Our research examines the role of one influential intermediary actor (the ISP), and will hopefully shed some light on the practical options.

1.6 Deliverables

The main deliverable of this research project will be thesis consisting of 6 chapters.

This proposal will be used for the first chapter. In the second chapter, a review of the current literature on Internet security will be provided, and we shall reflect on the state of the art in economics of information security. The chapter will start with a look at the cyber-security landscape, and gradually shift its focus on to botnets, and how they might be mitigated. The chapter ends by recognizing a research gap, and developing research questions and a conceptual framework. This will be accomplished by applying economic concepts to the underlying technical problem.

In the third chapter, the methodology that will be used to carry out the research will be presented, and the sources of data used for compiling a usable dataset introduced. The chapter will finish by constructing a measurement model and empirical hypothesis, based on the conceptual framework. The fourth chapter complements the third chapter by looking at a number of special issues involved in preparing the data.

This leads us to the fifth chapter, were we will have finally reached the stage of data analysis. Different sets of statistical instruments will be employed, and the results interpreted. The thesis concludes with chapter six, which looks at the policy implications of our findings, then reviews the limitations of the work, and provides suggestions for further research.

CHAPTER 2 – A REVIEW OF THE LITERATURE

2.1 THE CYBER-SECURITY LANDSCAPE

2.1.1 ONLINE SECURITY THREATS FACED BY END-USERS

During the last decade, the Internet has changed into one of the critical infrastructures in developed countries, to an extent that major disturbances in the Internet will significantly affect everyday life and business. This makes the issue of securing the Internet and dealing with cyber-security threats ever more important.

There are various levels of threat that the Internet faces: those targeting the infrastructure, governments, big corporations, SMEs and end users. In this section we will briefly list the major online threats that end-users face. The reason for focusing on end-users is that they are currently one of the 'weakest' links in cyber-security many of them (like grandma) don't have the necessary technical skills to defend themselves, and there are millions and millions of them, making the average Internet user rather easy targets (adapted from OECD 2007). There is another interesting reason why to focus on end-users: *botnets.* Botnets, which will be introduced in this section and elaborated on in the next, are armies of comprised end-user machines, that cyber-criminals exploit to amplify their power when performing their attacks. Basically put, end-user security has huge implications for the health of the Internet as a whole.

We shall start this section by making an inventory of the various threats end-users face.

VIRUSES, WORMS AND TROJANS

These names are without doubt the most famous of all security threats, as they were around even before the Internet took off. '***Viruses***' and '***worms***' are programs that replicate themselves without the user's permission or knowledge. They infect a host machine; execute a *payload*; and using various mechanisms (such as email, networks shares, USB sticks, etc) spread to other hosts. This is the general principle, but the payload differs exceedingly. It can show a harmless pop-up message for the fame of its creators; or it can perform more damaging acts, such as deleting user documents.

Although similar in actions, a technical difference exists between viruses and worms. Viruses attach themselves to other programs for replication, and hence to the user they seem to be performing the application's function. Worms on the other hand, are standalone programs with no useful functionality, and mostly replicate via network vulnerabilities. It should be noted however that these days these differences have become rather superfluous, due to the fact that the distinctions between worms and viruses (and other forms of malicious code) is blurring.

'***Trojans***' are another group of malicious software, which like the famous Greek Trojan horse, perform an apparently useful, harmless or even amusing function, while unknowingly to the user, also perform some malicious deed. One humorous example was a postcard application that displayed a group of dancing sheep on the screen,

while behind the scenes, it was in fact sending a copy of all passwords saved on the computer to a remote location.[1] The replication vehicle employed by Trojans is the deceived users themselves!

One use of trojans is to create a '***backdoor***' on the user's computer. This means planting a possibility in the system for the attacker to access it at will in the future. This could for instance happen by creating a new user for the attacker with administrative rights.

SPYWARE, KEY-LOGGERS, AND ROOTKITS

Nowadays, the trend for malicious software is to perform less noticeable damage (such as rendering documents useless), but rather to stay in stealth mode, and remain undetected for longer periods, in order to maximize the long-term (financial) gains of the attacker. One example is spyware. '***Spyware***' is software that "spies" on the user, by tracking her behaviour on the computer. For instance, it could send the sites the user visits to a remote location. This information can then be used for showing targeted, unsolicited advertising to the user (coupled with adware). Other example could include harvesting all email addresses the user corresponds with. They could even be deployed by governments for spying on their citizens, like the recent discovery of such acts by the Chinese government on the Tibetan exile government (UToronto 2009). One feature of spyware is that it's usually very hard to remove it once installed.

'***Key-loggers***' are a special category of spyware that sit in the background and capture all key strokes on a system, and send these to the attacker (or save them in a file for later retrieval). The most common use is for stealing passwords and capturing credit card information. When malicious software is designed to hide very deeply in a compromised system, it is named a '***rootkit***'. Detecting the existence of rootkits is very hard, as they obscure their presence from ordinary OS security mechanisms. Rootkits typically use extensibility hooks in modern operating systems.[2]

MALWARE AND BOTS

Security experts have adopted the generic term '***malware***' to refer to the many categories of malicious software[3] in existence these days, especially since the line distinguishing them is blurring. For instance, a malicious program might be a trojan, a worm, and a key-logger; That is, it is able to replicate on its own (via some vulnerability), but also fools users into distributing it, and when run, installs a key-logger on the victim's machine.

Malware has many criminal uses, but it is created and spread for financial profit (as compared to fame and glory, which was a bigger motive in the early 2000s). When malware infects a system, not only does it cause trouble for the computer owner, but also effects other computers connected to the Internet negatively. In economic terms, it creates a negative *externality*[4]. This is because many infected systems today are turned into a '***bot***' (or '*zombie*'), that is used by the perpetrators to:

- send out spam;
- launch Distributed Denial of Service attacks against other systems;
- conceal the tracks of hackers when they attack businesses;
- and many other nasty things

[1] A common trojan these days is the "rogue" anti-virus (malicious code pretending to be anti-virus software)

[2] When software containing the rootkit is executed for the first time, the rootkit integrates itself with the OS - as a device driver, explorer extension, etc. Traditional anti-virus and anti-spyware software have great difficulty scanning these areas.

[3] By malicious software we mean software that has a malicious intent unknown to its user.

[4] Externality is an economic term that can be loosely defined as side-effects of a transaction on parties not directly involved.

We will revisit *botnets* (armies of bots) extensively in section 2.2. At this point, it suffices to iterate the point that that malware targeting end-users is a serious threat for all those online, not just the end-users. Statistics point out that the problem of malware is getting worse as time goes on.

Traffic Interception

'***Traffic interception***' or '*packet-sniffing*' is a security threat that occurs at the network level. When performed illegally, it is a 'passive' attack, and can go unnoticed by the network users, security software, and even network administrators. Since a large amount of web, email, and voice traffic passes through the Internet unencrypted, an attacker that has access to the transmission path, and successfully intercepts data packets, can inspect them, and record any sensitive information for later abuse. Pulling off traffic interception is not always hard: if you use an unsecure WiFi hotspot, your traffic is interceptable by those in your proximity. Internet service providers can also, by their nature, see all your traffic. End-users can avoid the interception threat by sending sensitive information only via encrypted protocols (such as HTTPS).

Phishing and Poisoned Websites

A different category of Internet threats rely on deceiving the user in fraudulent ways. These threats fall under a general category of attacks known as *'social engineering'*. One of the most notorious of these techniques is '***phishing***'. In phishing attacks, the victim is presented with a web-page that looks identical to a legitimate online service that she normally uses, and she is hence lured into entering credentials and other sensitive information on the fake site. [1] The attacker records this sensitive data for fraudulent purposes. Figure 4 and Figure 5 show two examples of phishing, one targeting 'MSN', and the other 'Click and Buy'.

Mitigating phishing has two elements. The first element is to make users aware of this threat, and ask them to do certain checks (e.g., the URL is correct, the connection is secure, etc), before entering sensitive information on a website. The second element is that the targeted institutions need to actively find and takedown such sites. Here phishing becomes tied to malware: many of these fake websites are hosted on bots; hence the attackers easily move the sites to another bot once it is taken down, and their identities remain concealed as well.

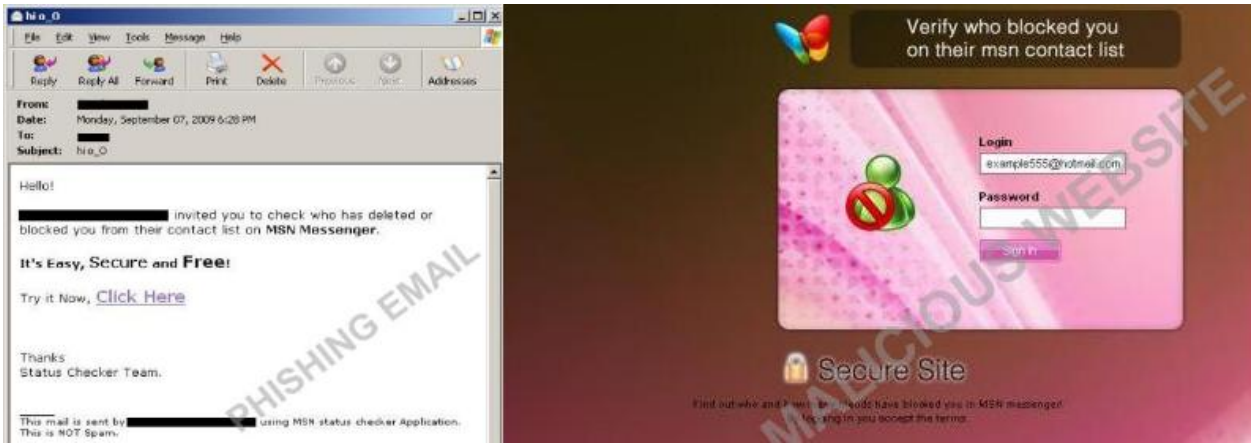

Figure 4 - 'See who blocked you on MSN' phishing attack (TrendLabs 2009b)

[1] The exact mechanism that the user is *'phished'* varies. Common tactics include sending a link via email (pretending to be from the user's bank and asking them to logon and check something); or registering a domain with misspelling of an actual domain, for instance, *ab**m**amro.nl*; or a more advanced mechanism like changing a user's DNS records (Anderson 2008) .

Figure 5 - Internet payment site Click and Buy phished; right: legitimate website; Left: phishing website (TrendLabs 2009a)

Somewhat related to phishing, '***poisoned websites***' are another security gaining popularity. In this attack, initially a legitimate website (be it a business, a non-profit-site), is infiltrated with malware, usually in an automated manner. Future users visiting the site now unsuspectingly face the risk of being injected with malware themselves; or having sensitive information stolen; and in the most harmless scenario, being bombarded by ads. This infiltration can often go undetected by the site owners for some time.

SPAM

'***Spam***' emails are often categorized as security threats, despite the fact that the majority of them can be considered as annoyances. This is not without reason: spam messages are frequently used to distribute malware and phish sites. Spam can be defined as *"an electronic message that is unsolicited and bulk (UBE)"* (Schryen 2007; Spamhaus). Unsolicited means that the recipient has not agreed to receive the message in advance. Bulk means that an identical message is sent to a large number of recipients. (A third condition, that the nature of the message must be commercial, is sometimes also added. In this case spam would be abbreviated to UCE, as compared to UBE.) [1] Spam can have multiple categories, as listed in Table 1.

Table 1 - Spam categories (adapted from Schryen 2007)

Spam category	Damage	Notes
Commercial advertising (UCE)	Nuisance[2]	This category includes sale of both legitimate and counterfeit goods. A famous example of a counterfeit ad campaign is the 'Canadian Pharmacy' selling pills.
Non-commercial advertising	Nuisance	Examples include political, and religious ideas
Hoaxes and chain emails	Nuisance	Trick people into believing something false (e.g., Nostradamus' predictions regarding 9/11), coupled with a recommendation to forward the email
Joe Job	Defamation	A forged email, apparently sent by a certain individual, where as the real goal is to damage the reputation of that person
Malware	Infection	-
Phishing fraud	Fraud	-
419 scams	Fraud	Also known as Nigerian spam, the sender claims to be a bureaucrat, banker or royal toadies, who want to cut you in on the financial deal of a lifetime, with the requirement that you pay some fee in advance.
Other scams	Fraud	Examples include pump & dump stock schemes, and pyramid schemes

[1] As can be guessed, all the terms in the definition are open to interpretations, so quite often there is no consensus on what constitutes of spam and what not!

[2] Nuisance still equals lost productivity and financial loss

It is worth mentioning that although spam has been ruled illegal in many countries during the past few years its volume has been constantly increasing, mostly due to botnets.

A MINI HISTORY OF MALWARE AND SPAM

Table 2 provides a history of malicious code, based on several sources. The most notable trends are that although mischievous code was written since the start of the computing era, until the early 1990s, the work mostly consisted of 'ethical hacking', that is the intent of the code writers was not too cause harm, but rather to learn and to provide proof of concept for ideas. In the 1990s, as malicious code turned dangerous, the antivirus software industry was born. Experts predicted that with migration to 32 bit operating systems, the AV industry would go out of business. The reality is however that this never happened. The Internet saved the industry, especially after several large-scale worm outbreaks in the early 2000s, which grabbed lots of media attention. Since 2004, malicious code has become more stealth, and is created mostly by organized cyber-criminal gangs with financial motivation.

Table 2 - Significant events in the History of malware (adapted from Anderson 2008 ch. 21; Goodman, Cormack, and Heckerman 2007)

Date	Event	
1960s	Students write computer games that if run as root, would give them priority access to shared computer time.	
1978	Shoch & Hupp write a program called 'worm' that replicates across the network and gives tasks to idle processes, and publish a paper about it.	
1984	Viruses first appear in public, after the thesis work of Fred Cohen, on how code could propagate itself from one machine to another	
1986	US computer fraud & abuse act written	+
1988	First Internet virus written by Robert Morris, which exploited a number of vulnerabilities, and ended up clogging the net.	
Early 1990s	Rise of the Antivirus industry, due to the growing virus problem. Experts predicted that the problem would go away once the move to "proper" operating systems would happen (from DOS).	
1997	As email turns into a vital communication platform, 'spam' also starts turning manifesting as a problem	+
Late 1990s	Interpreted languages and macro languages made the virus problem worse, up to the point that in 2000 macro viruses accounted for almost all the viruses. The net effectively save the AV industry	
2000	"Love-bug", a self propagating worm, sends itself to contacts in the victim's address book, with the subject "I love you"! Many copy-cat worms appear.	
Early 2000s	Flash worms such as 'Code-red' and 'Slammer' appear. These worms propagate by actively scanning the network for machines vulnerable to exploits and taking over them. Within hours a large number of users are affected!	
Early 2000s	Rise of spyware and *adware*. (Adware is software that bombards users with ads). Funnily enough, one of the malware writers sues AV companies for blacklisting them!	*
2002	EU Directive on privacy and electronic communications passed, making spam and malware illegal	*
2003	US CAN-SPAM (control of non-solicited pornography and marketing) act passed	+
2004	Microsoft releases Windows XP SP2, narrowing the window for malicious code. Social engineering finds a more important role in network attacks.	***
2004	Virus writing transformed from an amateur activity to an organized criminal economy in information goods. The malware writer's goal becomes to recruit machines for selling in cash to 'botnet herders'.	
Late 2000s	Cybercriminals, like the mafia wishes to avoid attraction, and hence move toward "manually-controlled exploit campaigns" – more frequent attacks but limited in scope.	
2007	A large, hard to battle botnet, 'Storm', comes to light. Storm has over a million zombies.	
2008	The shutdown of the McColo hosting service, a hub for botnet command and control, reduces the levels of worldwide spam by nearly 80%. The happiness is unfortunately not long lived.	***
2009	A new large botnet, 'Conficker', grabs attention by recruiting several million bots and proving very tricky to kill	***

blank = Anderson; + = Goodman et al; * = other sources

Conclusion

In this section we overviewed the common security threats that end users face when using the Internet. We introduced terms such as malware, botnets, and social engineering. Malware is a catch-all name for all forms of malicious software. It is used to compromise computers and turn them into 'bots' controlled by the criminals. Nowadays, most forms of online threats, even the ones that employ social engineering, rely partly on botnets. Figure 6 shows the differences between categories of malware in terms of visibility and criminal intent. Finally, we hinted that in general, online threats are moving towards profit and away from fame. In the next section we will look at statistics to see the extent of the threats discussed in this section.

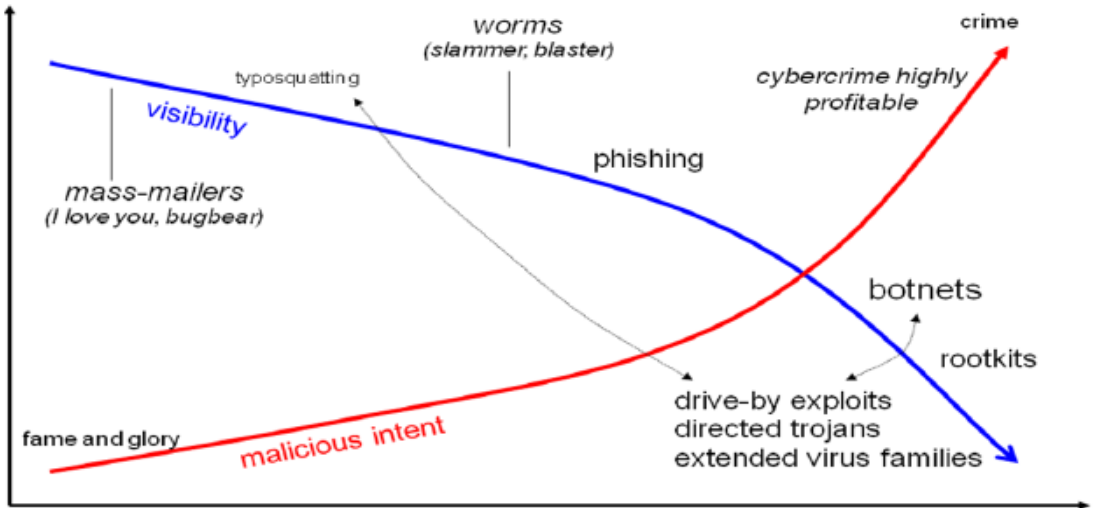

Figure 6 - Visibility of malware versus intent (OECD 2007)

2.1.2 The extent of the cyber-insecurity problem

Trends related to malware

How big of a problem is caused by cyber-criminals, and in particular malware? Estimates vary between experts, depending on the approach used to measure the effects, but in general, all sources put the economic loss caused by cyber-insecurity in the range of billions of euro. We shall investigate this matter by using the following sets of metrics:

- The growth levels of spam and malware;
- The percentage of individuals and businesses who fall victims to cyber-threats;
- The damage in monetary terms, especially from a social welfare perspective[1].

Spam growth

According to IronPort (2008b) and Cisco (2008) spam volumes constantly grew from 40 billion spam messages per day at the end of 2005 to an astonishing 200 billion spam messages per day in October 2008, a growth rate of 400% over 3 years (see Figure 7). At the end of 2008 the global spam levels dropped considerably due to the

[1] That is, by taking into account the costs of malware for society as a whole, and also the benefits it creates. Example costs include fraud and lost productivity, and benefits include profit for criminals and sale of security products.

shutdown of two botnet networks, however spam levels gradually grew back to their previous levels in 2009. If you consider the number of people globally online, it comes down to 12 spam emails per user, per day![1]

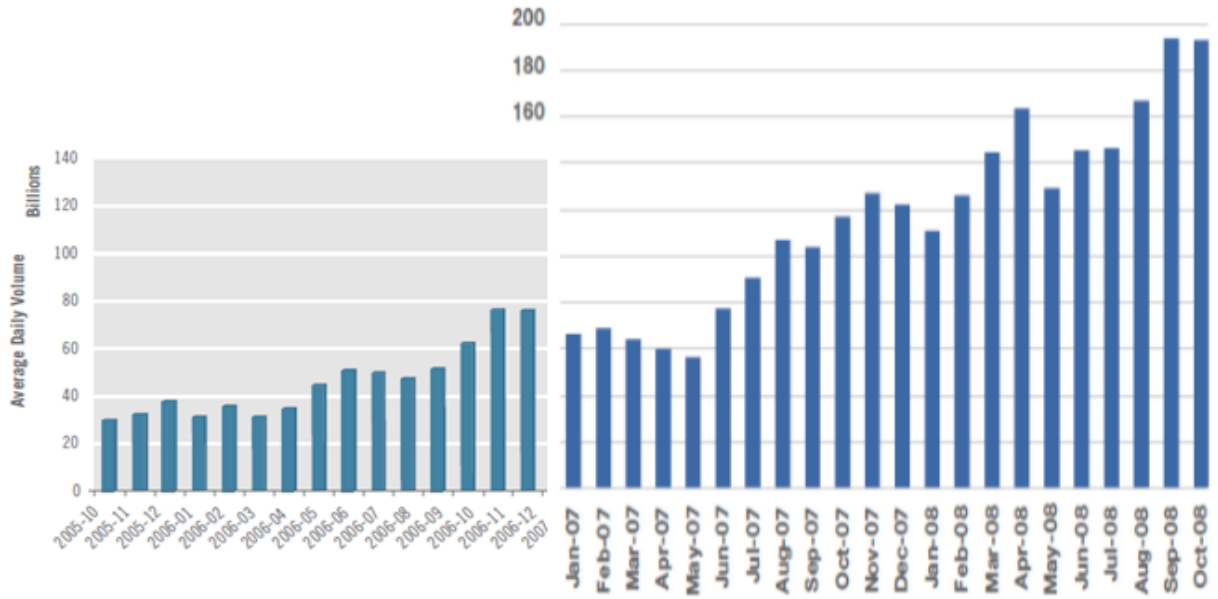

Figure 7 - Average daily spam volume - worldwide trends 2005-2008 (Cisco 2008; IronPort 2008b)

Other sources report similar magnitudes. TrendMicro (2009) gives the number of 115 billion spam messages per day at the end of 2008. MessageLabs (2008) reports that spam constitutes more than 75% of total email traffic , as shown in Figure 8.

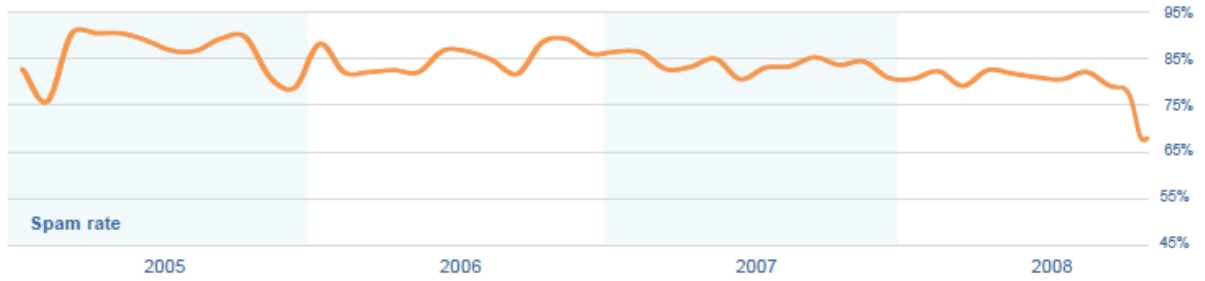

Figure 8 - Average global proportion of spam in email traffic (MessageLabs 2008)

Malware growth

The number of malware strains has grown exponentially over the past few years. Approximately 700,000 new malware are identified per month—an enormous increase compared to previous years (AV-Test GmBH as cited in TrendMicro 2009). This increase is largely due to auto-generated malware and causes significant headaches for security experts.[2] Enrique Salem, CEO of Symantec, announced comparable figures in this year's RSA conference: 10 million new malware signatures were discovered in 2008, equal to the number of malware created in the

[1] Luckily much of this spam is stopped by anti-spam technologies and never gets through to the user.

[2] Traditional anti-virus software works by scanning executable files for known virus patterns. With such high number of patterns to search for, not only is the AV signature list almost immediately outdated, but the software's scan engine also becomes extremely slow. Experts are looking into new approaches, such as using an application white-list.

previous 17 years combined! [1] These numbers mean that between one to two thousand new unique malware were created per hour in 2008.[2]

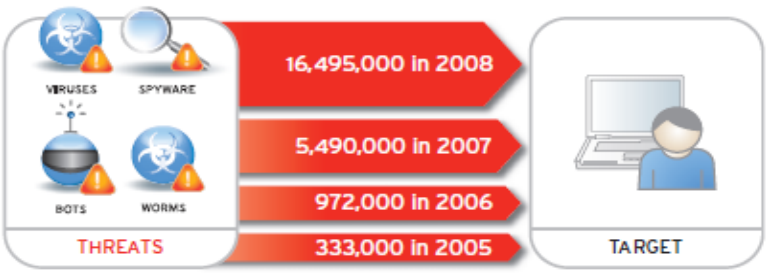

Figure 9 - New unique threats (TrendMicro 2009)

Poisoned websites

Provos and others (2007) performed an analysis on all pages indexed by Google in 2007 and found that one in ten web-pages were infected with malicious code, 70% of which were "legitimate" websites. IBM (2009) reports that only in Q4 2008, the number of new malicious websites surpassed the number seen in the entirety of 2007 by 50%; MessageLabs (2008) reports that the daily number of *new* websites containing malware rose from 1,068 in January 2008 to a peak of 5,424 in November; TrendMicro (2009) tells the story of an attack in May 2008 in which half a million websites were injected with a malicious script. All in all, it's rather easy for an ordinary user to stumble upon a malware-poisoned website.

Percentage of users infected with bot infections

Estimates for the number of computers recruited into a botnet through malware infections vary between security experts, but all of them put the figure into millions of machines. As an example consider the number of computers infected with Storm, the first botnet to catch media attention. Some researchers believe that as many as 50 million computers have been affected at some point in time. More conservative estimates put the number at 1.4 million computers infected and active in July 2007, with 900,000 new infections per month (IronPort 2008a). Even this number is truly extraordinarily, making the Storm botnet a powerful supercomputer in the hands of criminals. The situation has not improved since 2007. In February 2009, the Conficker/Downadup worm was believed to have affected as many as 12 million Windows computers, with Microsoft putting up a $250,000 bounty to find the person behind the malware (BBC 2009b).

In terms of percentages, Microsoft (2009) gives the lowest estimate, with a worldwide infection rate of 8.6 for every 1,000 PCs[3]. Although this figure is much lower than some dooms-day predictions, it still makes for an enormous number of PCs. One dooms-day predication was done by Vint Cerf of Google in 2007. He likened the spread of botnets to a "pandemic" and stated that up-to 25% of end-user machines could possibly be infected (BBC 2007). Panda Security also report a similarly high figure, stating that as of March 2008, 30% of computers on the Internet were infected with half of them being active bots (as cited in Bauer, Van Eeten, and Chattopadhyay 2008). Other vendors report figures in between these.

[1] Keynote webcast available online at http://media.omediaweb.com/rsa2009/webcast.htm?id=1_3

[2] Most of these malware strains belong to the same family and are variants of each other.

[3] This estimate is based on the execution of the Microsoft Software Removal Tool, an automated tool that runs once a month when automatic updates are enabled. Hence the results only include the most prominent malware strains targeted by the tool. Additionally, even among these strains, the numbers are biased downwards as the users that turn off automatic updates (and not accounted for in this estimate) are actually much more likely to be infected.

To sum up, we present Figure 10, an estimate of the combined number of bot infections per month (TrendMicro 2009). Note that if we consider other non-bot-related forms of malware as well, the numbers will be higher.

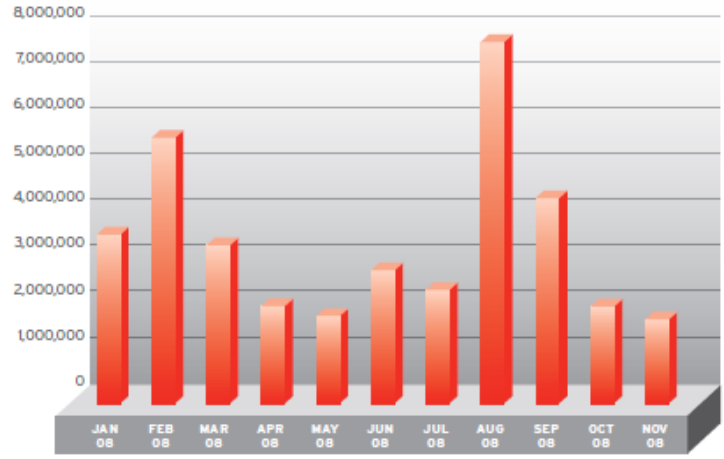

Figure 10 - Infections by bot related families (TrendMicro 2009)

Percentage of businesses affected by security incidents

The Computer Security Institute (CSI) conducts an annual survey among its members regarding security issues and practices. In 2008 about half of the members stated that they had experienced at least one security incident, causing an average loss of just under $300,000 per respondent (CSI 2008). The most expensive computer security incidents were those involving financial fraud, and the second most was dealing with "bot" computers. The respondents included US corporations, government agencies, financial institute, medical institutes, and universities.

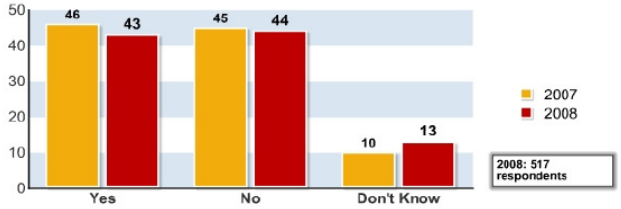

Figure 11 - Percentage of organizations that experienced a security incident(CSI 2008)

In another survey conducted for the British Department for Business Enterprise and Regulatory Reform, it was reported that 35% of all business encountered a security incident (BERR 2009). Both these surveys show that security incidents are common problem for all businesses.

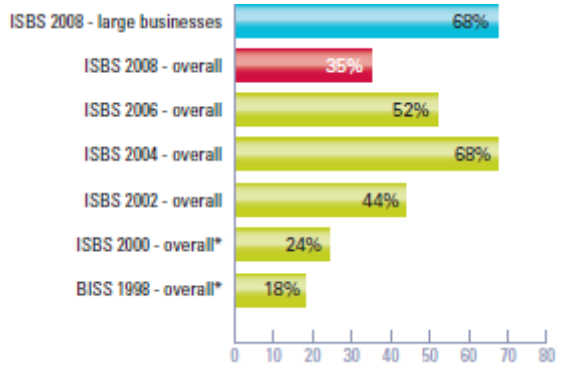

Figure 12 - Number of UK businesses that had a malicious security incident (BERR 2009)

Financial impacts of malware (and spam)

The previous section looked at the significance of malware and spam problem in terms of overall trends and percentage of people affected (rising and considerable!). To get a better sense of the scope of the problem, it would be interesting to measure the effects of malware in monetary value. The International Telecommunication Union has published a comprehensive report in this regards. The report measures the total welfare effects of malware, by considering both its costs and its revenues. A somewhat modified version of the financial categories used is presented in Table 3. The report also presents a conceptual model to quantify the legal & illegal cash flows created by malware (shown in Figure 13).

Table 3 - Financial impacts of malware (adapted from Bauer, Van Eeten, and Chattopadhyay 2008)

Costs/Revenues	Player	Types
Costs	Business-users (direct)	• Click fraud • Financial fraud (extortion, theft) • Other (ransom, loss of trade secrets, etc) • Payments to compensation end-users (if provided)
	Business-users (indirect)	• Cost of preventive measures (software, hardware, training) • Loss of productivity (e.g., time spent reading spam) • Loss of consumer trust • Wasted computational resources (e.g., bandwidth)
	Individual-users (direct)	• Financial fraud (including credit-card) • Credit-card and identity theft • Security software, computer repairs
	Society (indirect)	• Law enforcement • Opportunity costs of slower adoption of ICT • Attacks on the Internet infrastructure (born by all)
Revenues	Security Providers	• Sales of security devices, software, and services
	Other ICT firms	• Sales of hardware/software to criminals, users, and businesses
	Cybercriminals	• Middle-men: renting bots, selling tools, emails, identities, etc • Financial gain (fraud and theft)

Despite the authors disclaimer that due to the incomplete availability of numbers and their variability, it would be premature to determine the financial effects of malware with satisfactory reliability (Bauer, Van Eeten, and Chattopadhyay 2008), they offer a "patchwork of numbers", that gives a good indication of the extent of the malware problem:

- Costs of cybercrime (direct and indirect) for the US in 2005: $67.2 bn (FBI)
- Global cost of spam in 2007: $100 bn , a rise from the 2005 value of $50 bn (Ferris Research)
- Click fraud[1] in 2007: $1 bn (Click Forensics)
- Identity theft and other fraud (partially online) in the US in 2006: $49.3 bn (Javelin Strategy & Research)
- Costs of malware (viruses, spyware, and phishing) for US consumers in 2007: $7.1 bn (Consumer Reports)

- Combined global revenue of all security service providers in 2007: $7.5 bn
- Size of the global malware economy (legal and illegal transactions): $105 bn (MessageLabs)

[1] Click fraud is a mechanism for money extortion by fraudsters, in which they use misleading links (sent via spam), or automated bots to produce fake clicks on online-ads. This activity costs the advertiser money and gains money for the criminals, as the common practice in online advertising these days (for instance with Google AdSense) is for Google to charge advertisers only when a user has clicks on the ads; and part of the fee is paid as commission to the referring website.

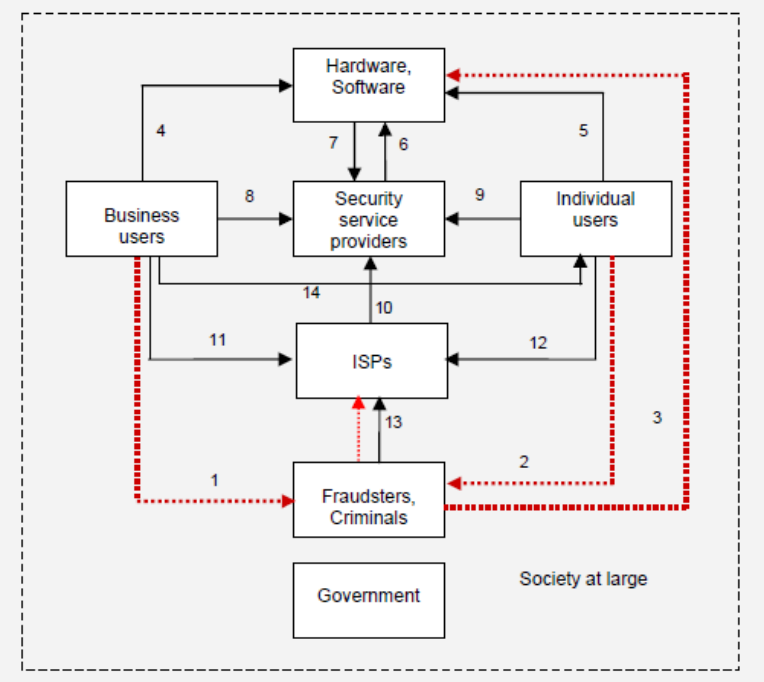

Figure 13 - Legal and potentially illegal financial flows related to malware (Bauer, Van Eeten, and Chattopadhyay 2008)

The authors sum these values to come up with a *conservative* estimate for the financial impact of malware and spam of 0.3% of global GDP, noting that if costs of slower migration to productivity-enhancing-ICT applications are added, the total impact will be even higher (Bauer, Van Eeten, and Chattopadhyay 2008)[1]. Non-conservative estimates can be much higher. The security firm McAfee, which like most security firms has an incentive to over-report cyber-crime costs (after all, scare tactics sell security solutions), recently reported a mind-blowing $1 trillion global damage caused by malware in 2008 (McAfee CEO Dave Dewalt in RSA-2009 conference[2]). Whichever figure you pick, cybercrime is costing the global economy and online citizens much money and a sense of urgency to fix it exists in the world.

Attacks on the government

The problems caused by malware are not limited to individuals or even businesses. There have been recent cases of infections and virus outbreaks in the public sector, and even the military. In one case the BBC (2009a) reported that machines inside the British and US governments had been taken over by malicious hackers from Ukraine.

Attacks on public infrastructures take three forms. The first form is attacks by cyber-criminals against the public websites, government PCs, etc, whether intentional or not, and for profit. By unintentional, we mean cases where the malware writer had no distinct goal of infiltrating the government or military, but rather, due to lax security, those systems were also hit while the malware was spreading.[3] Basically, the criminals use the same techniques and motives they use when conducting their business regarding corporations and other organizations.

[1] Gladly, the Internet in general contributes much more than this amount to the global GDP.

[2] Keynote webcast available online at http://media.omediaweb.com/rsa2009/webcast.htm?id=1_3 . The damage was calculated using Intellectual Property loss, and expenditures for repair.

[3] Reasons for directly targeting the government are like those involved in attacking other organizations, with the added value that the government has huge databases on citizens which might come in handy, and for some other specific forms of fraud.

The second form of attack comes from other governments. Depending on the scope of the event, they could be classified as either cyber-espionage or cyber-warfare. A recent and famous example is the case of the Chinese government's Ghostnet network (UToronto 2009).

A third form of attack, which should be differentiated from the previous two, is as a form of civil disobedience and online-activism. One such incident occurred after the June the 12th elections in Iran this year. During the two weeks of escalated conflict between the government and the people, members of the green movement performed coordinated DDoS attacks on government run news and media sites, as both a form of protest and to stop propaganda. [1]

Governments feel the heat to do something about cyber security. The FUD tactic mostly used is the threat of cyber-war, but in reality, the first form of attack (malware outbreaks) are much more devastating, as they are the most common, and influence the day-to-day continuity of public services

President Obama, himself a target of an online attack during his election campaign, has made cyber-security a top priority for the US. In April 2009, he ordered a 60 day review of the current situation, and ordered the creation of a new *'cyber tsar'* in his administration to oversee exactly this issue (BBC 2009c).

IMMUNITY OF THE CRIMINALS

A final way to look at the extent of cyber-crime is to look at an article published by the security firm Team Cymru (2006b) regarding "the cyber-underground economy" on the relative sense of immunity the cyber-criminals have. In this article, the authors use snippets of *public* conversations between cyber-criminals[2], to illustrate:

> *"the open arrogance the buyers, sellers, traders, and cashiers exhibit; the activities and alliances in which the underground denizens are involved; and the method by which they receive their ill-gotten goods; the blatant manner in which they advertise; and the personal data harvested every hour" (TeamCymru 2006b)*

These snippets are both amusing and illuminating. A few of them have been chosen and explained in Table 4. It is worth noting that although the article is reporting a problem for 2006, the mentioned IRC channels are still up and running! I decided to find the channels, and by spending just a few hours came upon a number of channels, for instance *#cc-power* and *#cc-master* in the *#undernet* server (CC stands for credit card). The criminals are still actively and publicly advertising their merchandise, requesting services, and openly trading with each other.

Team Cymru finish their article with the nasty truth that the cyber-criminals do not need to hide, or perform their transactions in private; they openly provide evidence of their crimes; and operate in a sense of immunity. The authors believe that the reason is a popular school of thought between decision makers that finding and prosecuting the criminals is too costly and resource intensive. In the next section, we will offer some more explanations for the status quo.

[1] The Distributed Denial of Service tactic was rather simple. Iranians across the world, especially from Europe, US (and even from inside of Iran), used a simple script to instruct their browsers to continuously fetch the homepage of the mentioned news sites (once every few seconds). The news sites could not keep up with the load of this seemingly legitimate traffic, and ceased working. The government's response was to shut all traffic from outside the country to its news sites!

[2] The snippets are captured from public IRC channels. Short for Internet Relay Chat, IRC is one of the oldest forms of real-time Internet text messaging, still in use with over 500 severs worldwide;

Table 4 - Buying, selling, and trading publicly in underground IRC channels (TeamCymru 2006b)

Conversation Snippet	Explanation
<A> Sell Cvv US(1$ each),Uk(2$ each)Cvv with SSN & DL(10$ each)and ePassporte Account with 560$ in acc(50$),Hacked Host(7$),Tut Scam CC Full in VP-ASP Shop(10$).shopadmin with 4100 order(200$), Tool Calculate Drive Licsence Number(10$).... I'm sleeping. MSG me and I will reply U as soon as I can !	Example of information goods offered for sale: credit card codes, social security numbers, and compromised hosts.
<A> Precautions 1. Use Fake Ip or Use a VPN While On This Server. 2. Do Not Use Your Real ID in picking any type of money 3. Dont give your real information to anyone unless you know him/her. 4. keep your self safe on [UNDERGROUND ECONOMY IRC NETWORK]. ThankS!!	Friendly advice on how to avoid getting caught!
<A> Sell cc's full info with PIN (debit, credit), COB's Laptops (alienware area51 = 500$, Dell inspiron 6100=400$, Scam pages (ebay, aol, paypal, egold, escrow, earthlink), track2gen (.exE) support 857 bins, 2000 bins (update bins), root. Payment (wu or e-gold).	Physical goods are also offered for sale (laptops). The payment in this case is via virtual-currencies.
<A> Name: Jason XXX <A> Address 1: XXX S University Blvd. <A> City: XXX <A> State: OK <A> Zip: XXXXX <A> Country: usa <A> Home Phone: (XXX) XXX-X991 Ext: <A> Date Of Birth: 12/8/19XX <A> Social Security Number: XXXX32199 <A> Mothers Maiden Name: Reaves <A> Drivers License Number: XXXX24766 <A> Drivers License State: OK <A> Secret Question: What is your pet's name? <A> Secret Question Answer: Joad <A> Name On Card: Jason XXX <A> Credit Card Number: 4492XXXXXXXX8831 <A> Credit Card Brand: Visa <A> Credit Card Type: Credit <A> EXP Date: 4/2006 <A> Credit Card PIN Number: <A> Card ID Number: X46 <A> Card Bank Name: OU Federal Credit Union <A> Card 1800 Number: 1800XXXXX9 <A> eBay User ID: XXX <A> eBay Password: XXXXXX	Identity theft – in this case the 'full info' of a person is provided for sale.
<B> I have Bank drops for Quick Cashout in(Hsbc,Wells, Lloyds, Citibank,Boa, Barclays,Woolwich,rbc) Contact me now for Fast Cash out..Deal is 50% each <D> Hello,I'm a professional MTCn confirmer if you have any order pending you can IM me,i have done so many transaction for different people and also i made different kind of transfer into account such as BAO, WELS,HSBC any body with full infos for this account who wanna transfer should IM me now and also i have BIN,EBAY SCAM PAGES,PHP bulk mailer if anyone is interested IM me all rippers keep off.NOTE I VERIFY FIRST.....	Extracting cash from the underground economy is the goal of many participants. Here you see advertisement of cashiers for both logical and physical account cleanups. Usually 50% goes to the cashier.
<A> i need who can confirmer westernunion female visa <B> speaking of wu, who can do females?	There are plenty of female miscreants too.
<A> I NEED DROPS FOR PHONES AND PDA's in Singapore Australia Austria Belgium Brunei Darussalam Canada China Denmark Finland France Germany Greece Hong Kong Indonesia India Ireland Israel Italy Japan Korea (South) Luxembourg Macau Malaysia Netherlands New Zealand Norway Portugal Saudi Arabia Spain Sweden Switzerland Taiwan Thailand United Arab Emirates United Kingdom United States	Drops are physical locations where stolen goods can be sent. The split is usually 50-50.
<A> JOIN #[CHANNEL] THE BEST HACKER CHANNEL!!! JOIN US ..!!! U CAN BECOME HACKER AND RICH...!!!!	Channel advertising!

Conclusion

In this section we looked at the extent of the cyber-insecurity using a variety of methods. These involved descriptive statistics, monetary value, government response, and criminal attitudes. All of them together paint a bleak picture. Experts do not believe that the situation will be auto-magically resolved any time soon (see Table 5 for 2009 forecasts).

Table 5 - Trend Micro (2009) Forecasts regarding Internet security threats in 2009

Prediction
Sophisticated blended threats are the new frontier
Social networking sites will grow as targets.
Alternative operating systems will be hit this year (Mac, Linux)
Mobile data is ripe for the picking
Microsoft—the eternal target—will continue its legacy of trouble in 2009
Social engineering will grow increasingly prevalent and cleverer
Broken DNS issues will continue to create headaches
Unlike the global economy, the underground economy will flourish
Identity theft will increase worldwide.
Spam volumes will continue to grow

To fix the problem, more research is required in order to increase our understanding of the cyber-insecurity problem, and to find possible remedies. The US National Science Foundation sponsored a think-tank to identify the greatest challenges that need to be solved in the 21st century. One of the 14 challenges: securing cyberspace (NAE 2008). The problem is serious and here to stay.

2.1.3 Reasons why cyber-security problems are hard to solve

Hopefully at this point it has become evident for the reader that the cyber-in-security is a serious problem. In this section, we shall take a brief look at reasons given in the literature for the current state of affairs.

Cyber-security is a wicked problem

Part of what makes cyber-security such a though issue to resolve can be tied to the fact that it is a "*wicked problem*". Rittel and Webber (1973) used the term *wicked problem* to describe problems that have no definitive formulation, are themselves symptoms of other problems, have no enumerable set of solutions, and the choice of explanation determines the nature of their resolution[1].

A good example is the case of spam. Putting aside the fact that there is still some disagreement to determining what qualifies as spam, there is much disagreement inside the technical community regarding every solution proposed to fix the issue. To illustrate this point, consider the following post in the security blog 'CircleID', which discusses the future of email authentication. (Lack of authentication in email systems is one of the loopholes that spammers have commonly abused; multiple solutions have been proposed in the last few years, such as SPF and DKIM). The author of the blog post is providing a discussion on the merits of the various solutions proposed:

> *"Over the years, many of us have collectively worked to provide a framework for authenticating email… Unfortunately somewhere on the path to protecting legitimate originators of email and the recipients of email, the larger community has gone astray…" [The author continues to provide*

[1] De Bruijn & Ten Heuvelhof (2000) identify two characteristics for wicked problems: disagreement on norms, and non-objectifiable information

> *his reasons and proposes a solution] "…There are many who would decry the changes outlined above. Some would claim that anonymity would be lost. This is simply not true…. Others argue that the risk of lost mail is too great…" (Hammer 2009)*

The point of this excerpt is to show the lack of consensus on concrete security measures, even on very specific measures. It goes without saying that in a dynamic environment where agreeing on any security measure requires much effort and is filled with political behaviour, solving larger security issues can be extremely difficult.

Path dependency and Internet governance

Part of the security landscape we see today is rooted in the fact that security mechanisms (such as authentication and encryption) were not built into the original Internet protocol suite, but are rather being added as an afterthought and in a somewhat ad-hoc manner (Anderson 2008). It could be argued that the Internet actually took off because of the relative simplicity of its protocol suite. Our goal however is not to judge whether this was good or bad, but to simply point out that a certain path dependency has been created, affecting much of the cyber-security landscape, and making its fixing very hard.

This path-dependency issue becomes more complicated considering how the Internet is run. The Internet is organized around *'autonomous systems*'- independently managed networks, most privately owned, or if public, managed at an agency level (Mueller 2009). There are currently around 30,000 "autonomous systems" active on the net (more details will be given in later chapters). Put simply, there is no central governing body, or forum (such as what we see in GSM) that can make final decisions or set standards on technical matters. Most decisions are done through consensus and "recommendations". This turns into a slow remediation process. We might be stuck with the current protocol inefficiencies (regarding security) for a long time.

Endless supply of software vulnerabilities

Many security threats rely on vulnerabilities (bugs) in the software running the Internet servers and clients. Despite much effort and discussion during the past decade, there seems to be no sign of a declining number of vulnerabilities (see Figure 14). It is true that the security of the operating system and core networking protocols has increased, and the number of vulnerabilities in these layers declined. But at the same time, attackers have shifted focus to the application layer (in tech-jargon, *"shifted up the protocol stack"*). Securing this layer is rather hard, due to the enormous number of applications in use (Scott Charney at the RSA Conference 2009; IBM 2009).

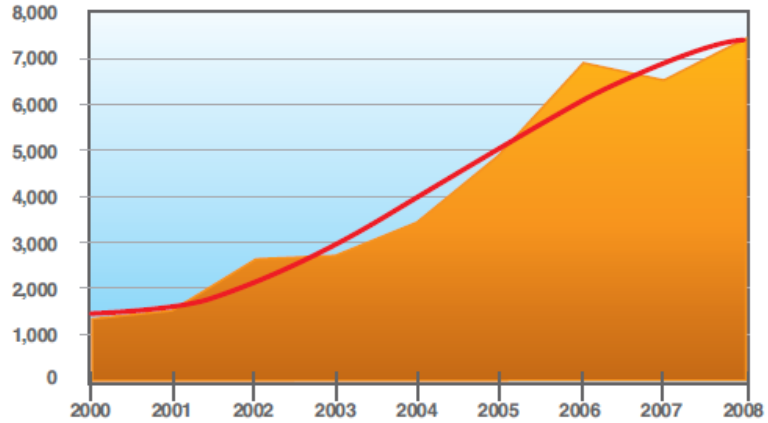

Figure 14 - Vulnerability disclosure, 2000-2008 (IBM 2009)

Interestingly enough, this lack of security in software is not unexpected. Shapiro and Varian (2000), and later Anderson (2008), present arguments that the software market does not reward security. The basis of the argument is that the value of software is largely determined by network externalities and the lock-in that it creates. In other words, the market rewards the dominating firm, which will be determined by the product that ships first, has the most number of bells and whistles, and provides ease of extensibility for other developers. These characteristics are all enemies of security. Microsoft's famous philosophy of the 90's – "ships it Tuesday and get it right by version 3" – is a reflection of this economic principle. This is the attitude the overall software market. Software vendors (and not just Microsoft) have incentives to eternally produce buggy code.[1]

INSUFFICIENT LAW ENFORCEMENT IN CYBERSPACE

Scott Charney, Microsoft's corporate vice-president of Trustworthy Computing, was the federal prosecutor for computer crime in the U.S. in the 1990s. Considering his background, it's not strange that he believes that "a percentage of the population will be always up to no good" (RSA Conference 2009) [2]. Society's answer to this population has been institutions such as police, courts, and others. Taking this into account, Charney states that cyberspace has several characteristics that make it even more attractive for the "up to no good" people:

- Access to large number of rich targets, due to the global reach of the Internet the Internet
- Possibility of automation of illegal activities
- Anonymity and lack of traceability of the criminals (due to both technical limitations, and limitations in cross-border law enforcement)

Team Cymru (2006a) provide similar arguments, pointing towards legal shortfalls, insufficient coordination, and lack of recognition of severity of cybercrime as some of the root causes of the cyber-crime epidemic. Anderson and others (2008a) also identify "fragmentation of legislation and law enforcement" as one of the main economic barriers to network and information security. Considering that these experts believe that the long-term solution for cyber-security requires changes in the mentality of law enforcement agencies and international law, it would be natural to conclude that the criminals are going to stay with us for some time.

MISALIGNED INCENTIVES

Another fascinating manner to look at the failure of security systems was put forwarded by Anderson (2001): information insecurity fails at least as much due to perverse '*incentives'* as it is due to technical failures.[3]

Anderson (2001) gives the example of ATM machine fraud in the U.S. as compared to Europe. In the U.S., if a customer disputes a transaction (e.g., complains about a withdrawal she did not make), the burden is on the bank to prove the customer is wrong; whilst in Europe the opposite is true - the customer have to prove their case to the bank. Of course, this is very hard for customers to do, which leaves the bank in a safer position. Interestingly enough however, after 20 years, the European banks ended up paying more for security and fraud. According to Anderson, the reason is that since the U.S. banks were liable for fraud, they had incentives to improve security, by for instance training staff to be more vigilant, or installing cameras in ATMs; while at the same time, their European counterparts were becoming careless.[4] A far reaching conclusion that Anderson draws from this story is

[1] The tendency of the software market to move towards dominant firms, creates 'lack of diversity', which is another security threat – it makes successful attacks more devastating. (Anderson et al. 2008a)

[2] The keynote webcast is available at: http://media.omediaweb.com/rsa2009/webcast.htm?id=1_4

[3] In economics, an incentive is any factor (financial or non-financial) that enables or motivates a particular course of action, or counts as a reason for preferring one choice to the alternatives. It is an expectation that encourages people to behave in a certain way.

[4] This is a case of "moral hazard" by the employees of the European banks.

that by assigning liability to the party who can '*manage the risks'* (the banks, versus the customers), the *incentive structure* changes in a way that benefits all parties in the long run.

We can use incentives to analyze the security outcome of the large online ecosystem. We see a certain degree of *misaligned incentives:* every actor is making rational decisions regarding her own security tradeoffs, but the sum of these decisions is a situation far from optimal[1]. Take the case of end-users: if they don't perceive much suffering in the case of a malware incident, they will not spend money on security software, or the time and effort necessary to educate themselves about security risks. This thinking is true for all actors. If an actor adopts a level of security that makes sense for her but harms others, we state that she is creating *'negative externalities'*[2].

The incentive perspective suggests that cyber-security can be fixed by changing and "aligning" the incentives of the actors in cyberspace. Some methods have already been proposed. For instance, Wash (2007) discusses ways of designing software that would provide incentivises to users induce better security choices; Loder, Van Alstyne and Wash (2004) put forward a mechanism called 'attention bonds', which target the incentives of the spammers, and argue that this would leave recipients better off than even a "perfect" filter that costs nothing and makes no mistakes.

Obviously, these are all very promising, making it worthwhile to better understand the underlying incentives of all actors. In this regards, Van Eeten and Bauer (2008) have conducted multiple interviews with Internet service providers, e-commerce companies, software vendors, registrars, and end-users, concerning how they make their security decisions, and the incentives they face. We will take a deeper look at these findings later in this chapter.

CONCLUDING REMARKS

In our review of the literature so far, we have had a look at the various online threats that end-users face, the severity of these threats according to various reports, and a list of root causes provided in the literature. The different views presented agree on two things. First, that the problem of cyber-insecurity is severe, and worsening. Second, the solution needs to be as much political, economical, and social, as it would be technical and engineering. These two make cyber-insecurity an interesting subject to research for a student studying Management of Technology.

Obviously, as the mentioned threats and perspectives to deal with the issue are quite large, some focus needs to be brought when researching this topic. We will continue our review of the literature by focusing on the most severe of all the threats, botnets; and we will solely make use of the incentives perspective to further analyze the situation and search for possible solutions, leaving out the other perspectives (such as discussions about Internet governance).

[1] This is similar to the famous "the tragedy of the commons" dilemma.

[2] In economics, an externality or spillover of an economic transaction is an impact on a party that is not directly involved in the transaction. In such a case, prices do not reflect the full costs or benefits in production or consumption of a product or service. An advantageous impact is called a positive externality, while a detrimental impact is called a negative externality.

2.2 BOTNETS IN DEPTH

In this section will take a more elaborate look at botnets, and investigate how they work, what uses criminals make of them, the business models that surround them, and some other related topic.

2.2.1 THE MECHANICS OF BOTNETS

Let us start by giving a clear definition of botnets. The Shadowserver Foundation[1] (2007a), defines a *botnet* as a collection of computers, connected to the internet, that interact to accomplish some (usually) illegal task. These computers are compromised, and are being used without their owner's knowledge. The compromised machines are called *drones* or *zombies*. The malicious software they run is referred to as *'bot'*.

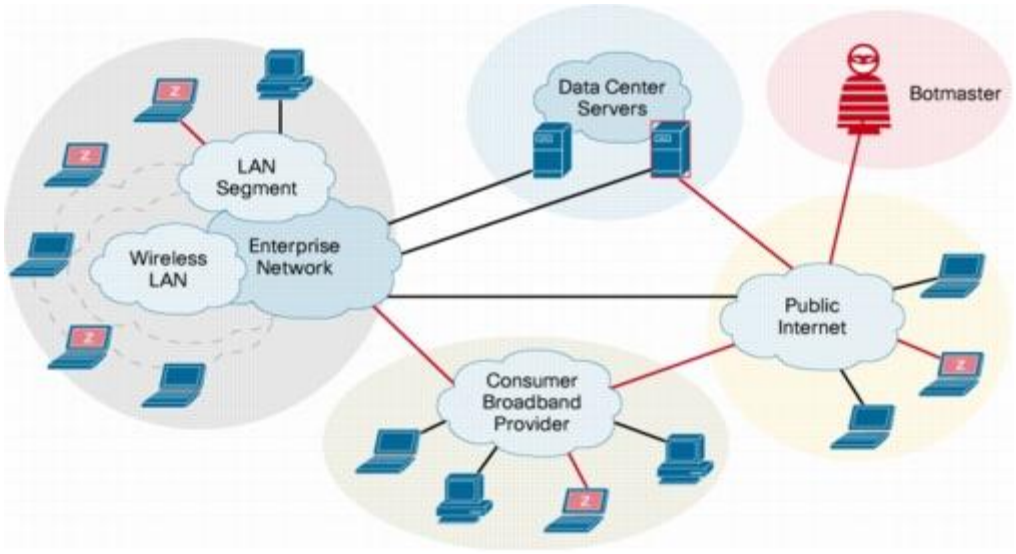

Figure 15 - A typical botnet with zombies (Cisco 2007)

BOTNET FORMATION AND PROPAGATION

It should be pretty obvious that the more hosts a botnet has, the more valuable and powerful it becomes. One goal of a *botnet owner* is to continuously grow his botnets by '*recruiting*' new bots (infecting more systems). The criminals are pretty creative and devious when it comes down to the techniques and opportunities they use for this. The infection mechanisms can be divided in two general methods:

- ***Automated infection***, through the exploit of software vulnerabilities or misconfiguration.[2] Bots contain a *scanning engine* that actively scans their surrounding IP addresses for vulnerable hosts. (The bots are behaving like worms).
- ***Infections requiring user interaction***: The user is deceived through various means to actually click and run malware. (The bots are behaving like trojans)

The specific methods used are numerous, and change quite frequently. For instance, there is the use of 'zero day bugs' – newly discovered software vulnerabilities for which the manufacturer has not yet released a patch; There

[1] The foundation, which is linked to by many sources, describes its activity on its website as gathering "intelligence on the darker side of the Internet", with the mission of "understanding and stopping high-stakes cybercrime in the information age".

[2] Vulnerabilities include newly discovered bugs, or older ones for which the user hasn't installed the patches. Examples of bad configuration are default passwords

are numerous delivery methods, such as spam emails, poisoned websites, USB sticks, etc; And the social engineering tactics exploit what's on people's minds at the specific moment, such as political races, the economic downturn, or a sports league. We will provide concrete examples when examining the Storm botnet shortly.

Typical uses of botnet

When describing the cyber-landscape, we named many of the illegal uses of malware and botnets, but for the sake of having them all in one place, we relist them in Table 6. It's important to know that many bots are written in a modular manner, meaning that the attacker can add new functionality to them as the need arises (OECD 2007). This is done by the bot downloading additional malware. Finally, botnets are *'self-sustaining',* as can be seen in the '*botnet lifecycle',* depicted in Figure 16.

Table 6 - Uses of botnets (adapted from OECD 2007; Shadowserver 2007a)

Usage
Locate and **infect** other information systems (botnet growth)
Steal sensitive **information** from each compromised system (sing key-logging, accessing the file-system, …)
Perform **D**istributed **D**enial **o**f **S**ervice attacks (the attacker orders all bots to bombard a target with traffic at the same time)
Send out **spam** (both for profit; and for distributing malware).
Host **phishing** sites (for fraud) and **mule-websites** with rotating IP addresses.
Engage in **click-fraud** (The bots click on ad-banners that earn the botnet owners money)
Host **warez** (pirated software), pornography, and other illegal content
Espionage (spying on infected users or intercepting network traffic)
Use as a **proxy/shell** for attacking larger information systems (removing IP trace)
Self **defence** (by killing anti-malware software, disabling updates, etc)

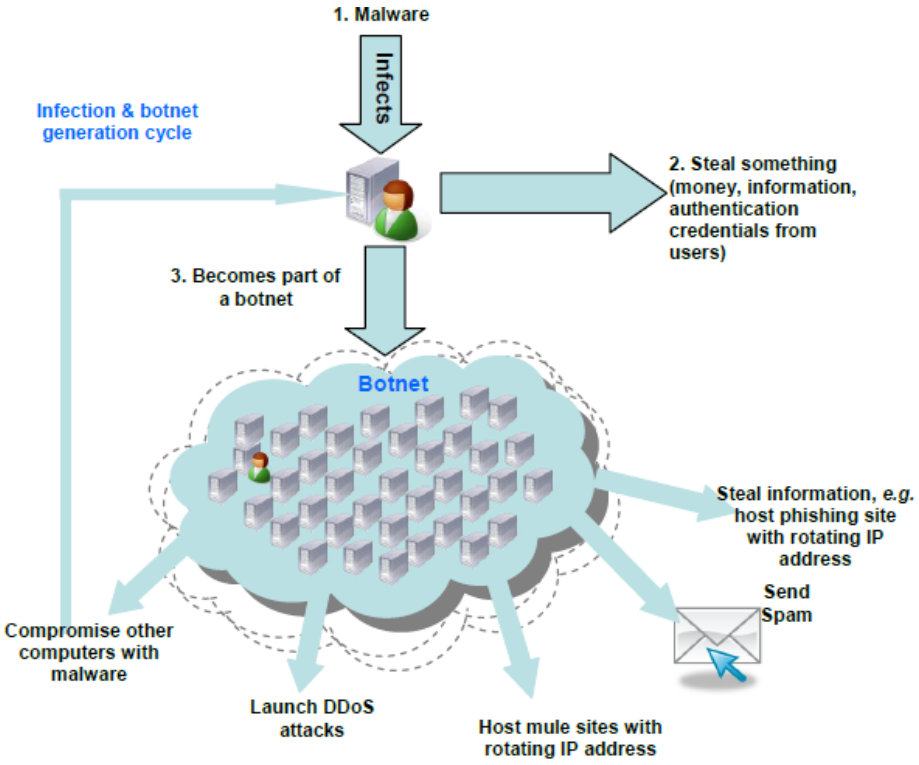

Figure 16 - The botnet lifecycle (OECD 2007; Shadowserver 2007a)

HOW BOTNETS ARE CONTROLLED

Since botnets are rather flexible and can be used for many purposes, they need to periodically receive commands from their *"masters" (*also known as *herders* or simply, *owners*). This is achieved via various *command & control mechanisms*.

In the first generation of botnets, the bots would all login to an Internet Relay Chat server, join a particular channel, and wait for their master to send them a message with new commands (Shadowserver 2007a). This approach was favoured for its simplicity but had several drawbacks. For instance, if a rival gang found the channel, they could hijack the bots![1] Similarly, authorities and ISPs could take down the botnet by blocking access to IRC.

Another C&C mechanism is to use an HTTP (web) server: the bots periodically access a webpage to receive their commands (Shadowserver 2007a). This method is stealthier then IRC, as the web requests get lost in the thousands and thousands of web-pages a user typically visits. However, if the server's domain is eventually found out, the same drawbacks exist. Figure 17 shows these mechanisms schematically.

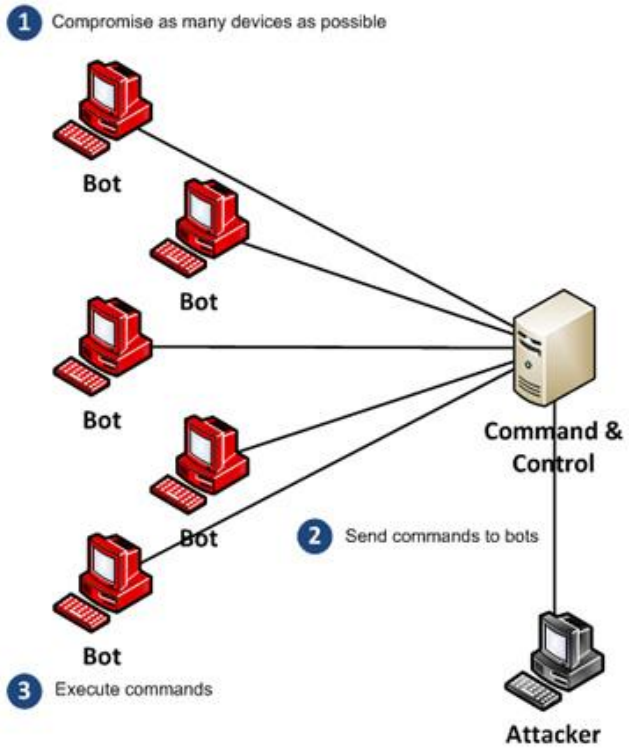

Figure 17 - Centralized Command & Control (image source: secureworks.com)

Other innovative mechanisms which are much harder to defeat, include peer to peer C&C structures (OECD 2007), and more recently, *random URL* based HTTP. In the latter, the web-domain of the C&C server changes every day. The bot generates a random list of domains everyday (e.g., 500 domains), and tries connecting to them all. The botnet herder only needs to actually register and setup one of these domains for each day to be able to send commands. The authorities however would need to monitor and close all of these domains *every day*, a very cumbersome endeavour.

[1] Encryption is used to mitigate this problem

2.2.2 Illustration of a botnet: the story of 'Storm'

In order to get a more concrete feeling of botnet propagation, uses, and behaviour, we will take a look at the ***Storm*** botnet. Storm was one of the first truly sophisticated botnets, and grabbed many headlines in 2007 and 2008. Storm received its name from the subject of the first email used to spread it: "230 dead as storm batters Europe" (Schneier 2007). Over its lifetime, the botnet owners came up with many enticing email subjects and web-pages, some of which are listed in Table 7. Exact figures on the size of the Storm botnet are hard to come by, with estimates ranging from 160,000 up to couple of millions (Blorge 2007). Storm had the usual botnet abilities, such as being a trojan, worm, bot, spam-engine, DDOS tool, etc, but in addition, it featured some interesting capabilities that made it hard to fight with. Schneier (2007) lists some of these as follows:

- Storm was designed as an 'ant colony' with segregation of duties (i.e., parts of the botnet do different tasks). Moreover, this colony uses a peer-to-peer C&C structure, making it hard to monitor and disable.
- Storm didn't create any noticeable performance impact on infected hosts, often shutting itself down for a while to avoid suspicion.
- Storm changed its distribution payload every 30 minutes, making it hard for anti-virus software to detect.
- Storm's delivery methods changed, including many websites and different email campaigns (Table 7)
- Storm actively protected itself against attempts to track or disable it. It performed 'counter attacks' against both computers scanning for it, and against the security researchers own sites, making sure "no one messes with it"!!

Of course, all this effort was not without financial motives. IronPort (2008a) reports that Storm made plenty of revenue for its owners, the biggest source of revenue coming from the ***"Canadian Pharmacy"***. Strom sent 1.5 billion spam messages daily, from 10% of its bots, promoting this pharmacy. The pharmacy used 100 new domains per day, and had 15 uniquely branded websites. The sites offered counterfeit-pharmaceuticals, (shipped from China and India), rather than brand-names from Canada. Buyers usually received the drugs, and the whole business, made possible by Storm, created an estimated revenue of over *$150 million* yearly! Table 7 lists some other revenue sources as well.

Table 7 - Social malware methods used for recruitment and revenue generation by Storm (IronPort 2008a)

Vehicle	Purpose	Explanation	Technique
Malicious "anti-spyware" sites	*Propagation*	Purport to offer a free scanner that will alert computer users to infections on their system. In reality, the user is downloading and installing malware.	*Social engineering via website*
Spoofed NFL site	*Propagation*	Active in the football season, promoted a NFL season game tracker application, which was in fact the Strom malware!	*Social engineering via web / email*
Spurious YouTube site	*Propagation*	Spam purporting to show a video clip featuring the recipient, would direct recipients to a site with the YouTube logo and a link to follow and hit "run" for the actual video. In reality, malware is downloaded.	*Social engineering via web / email*
Fraudulent e-cards	*Propagation*	Sent out on Valentine's Day and similar occasions, the messages announced an e-card from someone for the recipient. If the recipient clicks through to the website, they downloaded Storm malware.	*Social engineering / email spam*
Free Games, Psycho Kitty	*Propagation*	Targeted at younger demographics, these increasingly clever websites that look appealing and fun actually infect visitors' computers.	*Social engineering via website*
Vulnerabilities in widely used software apps	*Propagation*	Malware creators identify vulnerabilities in popular, legitimate software. They take advantage of a vulnerability by inserting active code which exploits the applications' flaws – and compromise the system.	*Software vulnerabilities*

Vehicle	Purpose	Explanation	Technique
Blog comment spam	*Propagation*	Spam comments on legitimate blogs include links leading to sites that infect computers.	*Web Spam*
Excel attachment spam	*Propagation*	Messages that included Excel file attachments. Used only for testing response rate and infiltrating anti-spam systems (not successful).	*Email Spam*
MP3 attachment spam	*Revenue*	Purported to be songs from well known artists, but instead the audio files contained ads that pushed stocks in "pump and dump" schemes!	*Email spam*
PDF attachment spam	*Revenue*	A tool for stock scams, they feature PDF attachments that look like well designed, legitimate investment newsletters.	*Email Spam*
Pharmaceutical spam	*Revenue*	The main source of revenue for the storm botnet, the spam messages direct recipients to credible-looking sites offering drugs like Viagra and Cialis for sale. The orders were filled with counterfeit pharmaceuticals.	*Email Spam*
Phishing spam	*Revenue*	Directed recipients to apparent financial management sites where their personal and financial information gets collected for criminal purposes.	*Email Spam*
Money mule spam	*Revenue*	Offered recipients work-from-home jobs transferring money through their bank or PayPal accounts for a commission	*Email spam*

Interesting enough, Storm vanished in Mid-2008 without a clear reason. Microsoft gives some of the credit to MSRT (Malicious Software Removal Tool), but it is also possible that the botnet has morphed into a bigger and scarier creature (TrendMicro 2009). To this day, Storms controllers have not been identified.

2.2.3 Organization of the botnet economy

The underground economy behaves like a highly specialized market, with segregation of tasks and services based on expertise (Bauer, Van Eeten, and Chattopadhyay 2008; OECD 2007). In the case of botnets, a simple value chain consists of the following actors (shown in Figure 18).

- *Crackers* (or malware creators), who engage in finding software vulnerabilities and writing code to exploit them. They sell their malware, sometimes in the form of 'do-it-yourself' kits;
- *Bot herders*, who acquire the malware (or malware kits), and spend most of their time distributing the malware, 'recruiting' new bots, and defending their network. They offer their botnet capacity for sale;
- *Fraudsters,* who rent bots from the herders, and use it to perform / launch / host their different schemes and attacks;

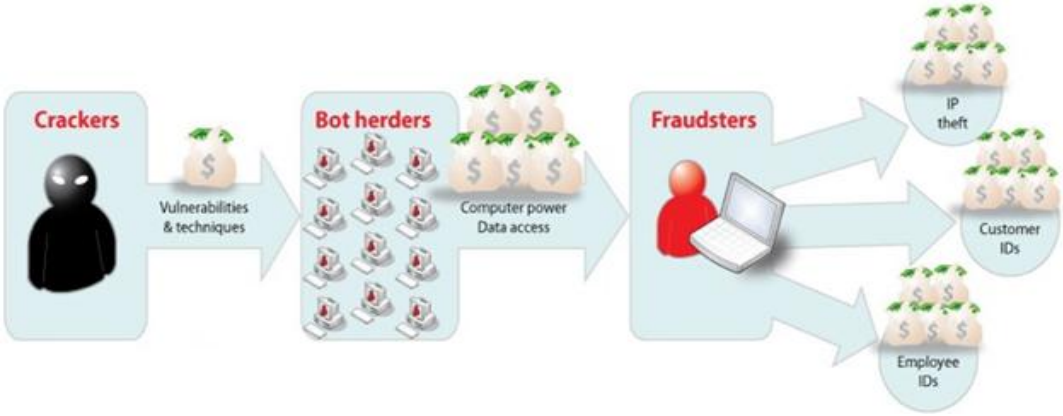

Figure 18 - Division of labour in the botnet value chain (image source: identitytheftblog.info)

The complete value chain contains many more actors. For instance, as we saw in section 2.1, *'cashing out'* is a critical part of the underground economy, undertaken by *'drops'.* Figure 19 shows a more elaborate picture of the division of labour.

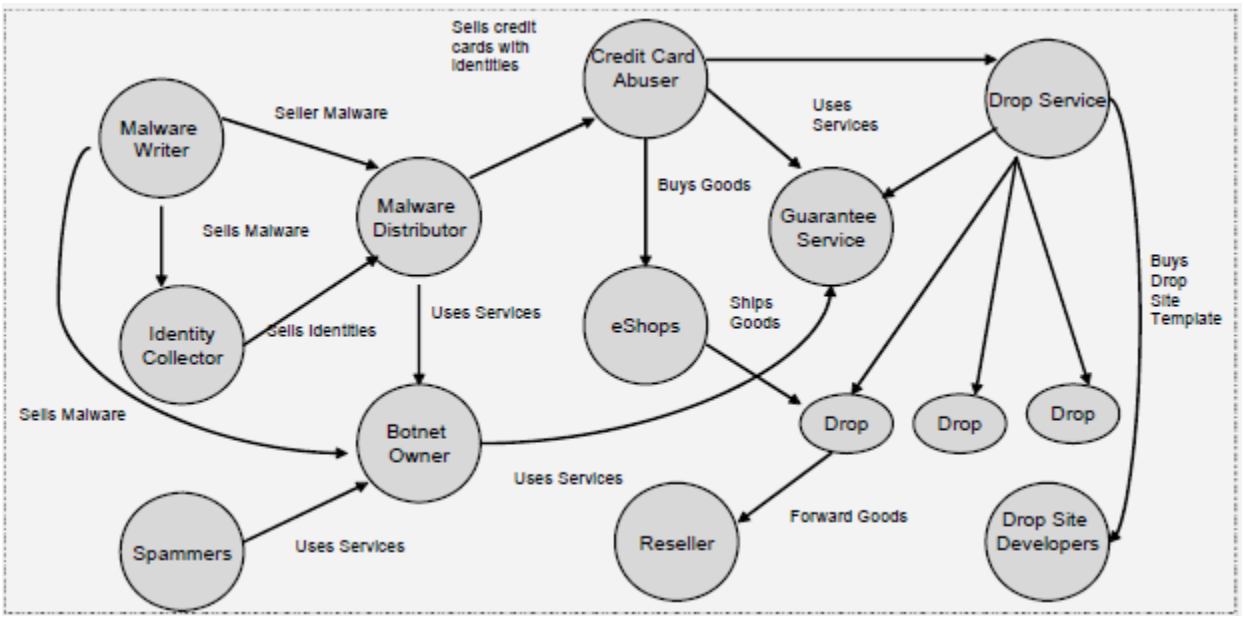

Figure 19 – Division of labour in the underground economy (MessageLabs 2007, as cited in Bauer, Van Eeten, and Chattopadhyay 2008)

There is much money to be made in the cyber-crime world. Consider this: gangs are offering computer programming graduates from Moscow's technical universities up to $7000 a month, were as the average salary for a professional Russian is $640 per month (TrendMicro 2009)! Of course, this value comes at the expense of the other actors. For instance, spammers make money, even though only 1 in every 12.5 million spam emails results into a sale (BBC 2008b), simply because the cost of sending spam is nearly zero. The poor recipients however need to spend considerable amounts of time deleting the spam. Criminals are causing severe negative externalities.

We will finish this section with a remark about factors that 'enable' botnets. Although botnets are the brainchild of smart cyber-criminals, there are various social and technological factors that enable their propagation and sustainability, e.g., they are economical, untraceable by the law, etc. Many of these were discussed in the previous sections. OECD (2007) lists two other technological enablers that we haven't yet discussed. The first factor is the growth of b*roadband* Internet. 'Always on' users with fast connections are ideal the targets for bot herders, for two reasons: they are faster to recruit, and more useful once recruited. A second factor is the increasing number of *services online,* which increases the number of targets for attack and exploit.

2.3 Botnet Mitigation

Having identified botnets as a serious cyber-threat (as they act as a platform for most other threats), we now turn our attention to ways this threat can be mitigated. The first answers that come to mind for botnet mitigation would be to target the criminals, or to increase end-user security vigilance. These are of course easier said than done, and in fact, the complete range of possible remedies is much broader. For instance, liability could be forced on software vendors, service providers, or other parties.

As briefly discussed in section 2.1, the status quo regarding Internet security, including the rise of botnets, is the result of various shortcomings, be they technical, legal, or economical. In this section, we will apply some basic economic concepts, such as *externalities*[1], *incentives*[1], etc, to further investigate the botnet problem. From an economic perspective, one path to mitigating botnets would be to change the incentive structure of some actors. Before we can propose any such steps, it is important to understand how actors are currently making their security decisions. Van Eeten and Bauer (2008) have performed an extensive set of interviews with market players regarding this matter, which we will review in this section.

2.3.1 Actor analysis

Table 8 lists the major actor groups participating in the cyber-security landscape. We will give a brief description of each actor; the strategies they have adopted; the incentives underlying this strategy; and finally, the externalities they create in the value net; (especially with regards to botnets and their cousin, malware).

Table 8 – Major actor groups affecting Internet security (parts adapted from Van Eeten and Bauer 2008; OECD 2007)

Actor Group	Notes / Sub-groups
Malicious actors	• Organized criminals • Fame seekers & copy-cats • Insiders
End-users	• Home users, • Small-Medium size businesses • Large end-users (government institutions, retailers, enterprises)
Software vendors	
Hardware manufactures	
Security service providers	e.g., Anti-virus firms
Internet governance bodies	• Consumer protection agencies • Internet bodies (ICANN, IETF, etc) • Industry associations (e.g., ETIS) • Regulatory agencies (national & international) • Policy making bodies (national & international)
Incident response	• CSIRTs (Computer Security Incident Response Teams) • Law enforcement
E-commerce companies	Including financial service providers
Domain Registrars	
Internet service providers	Including access providers, and hosting providers

According to Van Eeten and Bauer (2008), three situations can emerge regarding security externalities that actors create:

i. *The actor doesn't create an externality*; i.e., she bears the cost of protecting against security threats;

[1] See previous footnote for a definition of these words

ii. *The actor creates externalities which another actor capable of managing it absorbs;* i.e., she makes security decisions that deviate from the social optimum (due to lack of incentives or skills); However, another actor, capable of mitigating the size of the externality[1], willingly internalizes these costs;

iii. *The actor creates externalities that are imposed on society at large, or on unwilling actors.*

MALICIOUS ACTORS

Malicious actors can be grouped into five categories, based on their motivation and skills. These categories are:

- *Innovators,* who are individuals that for the love of the challenge, devote time to finding security holes;
- *Amateur fame seekers*, novice users who use ready-made tools to grab (media) attention;
- *Copy-cats*, hackers who desire celebrity status in cybercrime community by recreating simple attacks;
- *Insiders*, ex-employees and ex-contracts who want revenge, and do so by abusing their security privileges;
- C*riminals,* who are highly motivated, organized, extremely powerful, and are in the game for profit (OECD 2007).

The strategies adopted by the most vicious of the malicious actors, the criminals, are: extracting profit, running low by adopting stealth tactics; and defending their territory. This group of actors are the creators of the whole botnet problem, and for their personal profit, impose huge costs on society as a whole.

END USERS

End-users can be categorized into three groups: home users (and individuals); small to medium businesses; and large end users (public institutions and global corporations).

Home and SMB users create the most negative externalities among the legitimate actors. Part of the externalities is absorbed by other players, such as ISPs (Internet service providers) and FSPs (financial service providers). The negative externality is caused by risky online behaviour, in combination with not employing security software. This is outcome is a result of several facts: users don't understand how malware works, and often don't know when they become infected; they perceive the chances of being hit, or the harm that malware will do to them, as low[2]; they find paying for security software as unacceptable[3]; and lastly, due to bad design of security software, users do not like to install it.

It's important to know that things are all not bad for security: users are becoming more concerned and slowly adopting anti-virus and firewall software. (Van Eeten and Bauer 2008)

Large end users usually have dedicated IT staff that understand security risks, take precautionary measures, and handle incidents. Still, some believe that they are under investing and creating negative externalities – both in terms of malware infections, and more critically, in terms of "loss" of confidential citizen data. Incentives that effect security decisions of large users include: avoiding brand damage; legislation (both liability and compliance); and minimizing lost productivity.

Some disincentives of large users include: possibility of vulnerability patches "breaking" things; tradeoffs such as security versus availability, speed, or usability in their business; and finally, the monetary benefits of security measures not being very explicit. (Van Eeten and Bauer 2008)

[1] Meaning that this other actor can manage the risks through security measures

[2] e.g. thinking "I don't have anything important on this computer, so who cares if it becomes infected"

[3] e.g., Believing that it should be already included on the computer

In summary, end users create negative externalities, although not deliberately. Part of this is born by them and by other parties who see dealing with their externalities as a "cost of doing business".

SOFTWARE VENDORS

One could argue that half of the botnet propagation is rooted in exploitable software (with the other half being social engineering and user behaviour). We have previously discussed the role of software vulnerabilities in the whole cyber-insecurity problem, and some of the incentives that lead to it. Here we will re-examine these arguments. Please note that that software market is very diverse and are arguments are about the mainstream market, as in some specialized markets such as defence, security is a critical requirement from the start.

The most important disincentive is that the market does not reward security – at least before a firm becomes a dominant player. Rather, it rewards characteristics such as extra functionality, ease of use, and compatibility, all of which have tension with increased security. Developing secure software increasing vendor costs, while many times, it inhibits or distracts from compatibility and functionality. (Anderson and Moore 2007)

However, after a firm becomes dominant, two factors come into play that could act as incentives for increasing security. One incentive is avoiding (or mending) brand damage, similar to what happened to Microsoft in the early 2000s. After a series of spectacular worm attacks caused Microsoft's reputation to tarnish, it began taking security more seriously and started an internal code-review campaign in early 2002, which eventually resulted in the release of Windows XP Service Pack 2 (Thurrott 2004; Van Eeten and Bauer 2008). The other incentive is the cost of vulnerability patching. A patch consisting of 2 lines of code might take 3 months to test and release (in extreme cases), adding up to enormous costs.[1] To escape the costs of vulnerability patching, vendors need to invest more in security upfront. (Van Eeten and Bauer 2008)

An incentive mentioned in the literature that can work both ways is "user discretion" – that is users are responsible for their systems and the security decisions they make. This could be a disincentive for security, as it dumps liability. On the other hand, irresponsible user actions can still cause negative publicity for the vendor, so this argument goes only so far. (Van Eeten and Bauer 2008)

The net result: prior to market domination, strong incentives against security exist, but after domination, this changes. In either case, software vendors don't bear the full costs of software insecurity, and cause severe externalities. To be fair, the externality they cause is lower than the total cost of insecurity, as in a "perfect" market, users might actually choose software with a lower degree of security and remedy the problem with other countermeasures. Of course, lock-ins and information asymmetries deprive the perfect market.

SECURITY SERVICE VENDORS

This group of actors has a straight-forward incentive of selling its security solutions, and are in fact benefiting from the botnet problem. This creates a tendency for them to *over-report* losses, disasters, etc, and deploy scare tactics to promote their solutions. A recent and somewhat amusing example is a report by McAfee (2009) calculating the 'carbon footprint' of spam. The drive to sell security solutions often leads these vendors to promote the solution to malware as being mostly technological.

HARDWARE VENDORS

Hardware vendors (PCs & network equipment) do not have a direct role in the botnet / malware problem. They do have an incentive for growing sales, which will to some extent benefit from trust in the Internet and its usefulness.

[1] Installing patches is also costly, especially for the enterprise customer, so they might decide to ignore some of them, which has the possibility of costing negative publicity for the software vendor.

This is not a major concern however, as although some people might shy away from the Internet due to security fears, most citizens and businesses are not troubled, as evidenced by the continuous growth in the number of Internet subscribers and broadband penetration worldwide (source: ITU and World-bank statistics).

For PC vendors, there is a story that some customers expect their computers be safe 'out of the box' [***]. This means that the hardware vendors would need to bundle security software (e.g. anti-virus) on their machines. Such a strategy is beneficial for the hardware vendors as they gain commission on such sales.

INTERNET GOVERNANCE BODIES

Regulatory agencies and policy making bodies: Delving deeply into the incentives that governments have in solving cyber-security is not part of this thesis, but suffice to say that for a variety of reasons, such as fostering economic growth, and protecting their own systems online, they do wish to do something about cyber-insecurity. Anderson and Moore (2007) review some of the options available to governments.

First, in the mid 1980s, the US government (and later on NATO) tried to solve the 'lemons market' problem of security with the introduction of *certification*, and criteria schemes (the "Orange Book", and the "Common Criteria"). Both attempts were unsuccessful, due to among others, *adverse selection* (Anderson and Moore 2007). A second strategy has been to pass *regulation*, such as the HIPAA and Sarbanes-Oxley act in the US. Although these laws might have some effectiveness, they put disproportionate burden on small to medium sized businesses, and distort security markets (Anderson and Moore 2007). Direct regulation targeting unsolicited bulk email, computer intrusion, etc., also exist, but obviously are not having much deterrence on the cybercriminals. In short, the results of regulation are doubtful and speculative. Finally, s*elf-regulation* has also seen mixed results, working in some cases, such as patch-management, and failing in others , such website approval seal (Anderson and Moore 2007).

It seems that governments are aware of the limitations of current approaches and are searching for better policy options. This has fuelled interest in academic research in cyber-security. In a report, commissioned by the European Network and Information Security Agency (ENISA), Anderson and his colleagues (2008b) discuss some of the practical options available to EU governments for improving security failures. Their recommendations are listed in Table 9. Other sources put forward other ideas, such as awareness campaigns, international treaties, etc.

Table 9 - Policy recommendations to amend market failures for cyber-security within the EU (Anderson et al. 2008b)

Recommendation	
1.	The EU introduce a comprehensive security breach notification law
2.	The EC or Central Bank regulate to ensure publication of robust loss statistics for electronic crime
3.	ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs
4.	EU introduce a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines, with a right for users to have disconnected machines reconnected by assuming liability.
5.	The EU develop and enforce standards for network-connected equipment to be secure by default
6.	The EU adopt a combination of early responsible vulnerability disclosure and vendor liability for unpatched software
7.	Security patches be offered for free, and patches be kept separate from feature updates
8.	The EU harmonise procedures for resolution of disputes between customers and FSPs over electronic transactions
9.	The EC prepare a proposal for a Directive for proportionate and effective sanctions against abusive online marketers
10.	Research, coordinated by multiple stakeholders, to study what changes are needed to consumer-protection law…
11.	ENISA advise the competition authorities whenever diversity has security implications
12.	ENISA sponsor research to better understand the effects of IXP failures; and insist on best practices in IXP peering
13.	The EC put pressure on the 15 Member States that have yet to ratify the Cybercrime Convention
14.	Establishment of an EU-wide body for facilitating international cooperation on-cyber crime, using NATO as a model
15.	Regulations introduced for other purposes should not inadvertently harm security researchers and firms.

Industry associations, wishing to avoid government intervention, come up with various 'best practice recommendations' regarding Internet security, to promote a form of self-regulation in the ISP sector. We shall take a look at these recommendations in an upcoming section.

Consumer protection agencies play a somewhat double role in terms of positive and negative effects on cyber-security. On the one hand, these groups push for liability laws, and protection of customers against fraud, which is good for security. On the other hand, they pursue privacy laws and raise alarms about ISPs spying on end-user communications, which could inadvertently work against security (as discussed under the ISP actor).

FINANCIAL SERVICE PROVIDERS

Financial service providers (such as banks), and more generally, e-commerce companies, are unique among the actors in the manner that they *internalize* part of the damage caused by malware and botnets, by compensating end-users for online fraud. The reason for this strategy is clear: the enormous benefits that increased online transactions has for them.[1] (Van Eeten and Bauer 2008)

In order to achieve increased transaction volume, FSPs priorities are to keep their systems user friendly, usable, and available, even at the expense of security. Trust is also critical in financial transactions, and FSPs maintain trust by compensating users for any problems that might be caused by insecurities. A good example is the credit-card system: the system lacks 2-factor-authentication used in debit cards. This makes it very easy to shop with your credit-card, but it also increases cases of fraud ('card not present' attacks). The fraud however is not a deterrent for online-shoppers, as disputed transactions can instantly be revoked (first the money is refunded, checks are done afterwards). The results are astonishing: 358% growth in online and telephone transactions between 2001 to 2006, compared with122% growth in fraud during the same period. The global fraud rate for VISA in 2006 was 0.051%. (Van Eeten and Bauer 2008)

The process of internalizing costs (which the FSPs do for their benefit) is *socially optimum*, as these actors are in a position to manage the risks (Van Eeten and Bauer 2008). An example of being in a position to manage the risks is that credit-card companies are very good at detecting unusual purchasing patterns indicative of fraud, and might stop a fraudulent transaction even before the card-owner notices .[2]

DOMAIN REGISTRARS

A domain name registrar is an organization or commercial entity, accredited by ICANN[3] (or its delegates), to manage the reservation of Internet domain names for the public, in accordance with specific guidelines (Wikipedia 2009b). In simpler words: if you wish to have your own domain, you buy it from of the registrars. But how are registrars connected to Internet security and specifically botnets? One part of the answer lies in a technique used for hosting phishing sites, command & control servers, and malware delivery sites, known as *fast-flux.*

To understand fast-flux we must briefly explain how the DNS system works. Every host connected to the Internet has an 'IP address', which is used when other systems wish to communicate with the system. Now, when you are accessing a website, such as www.tudelft.nl, your computer is actually sending requests to a web-server with the

[1] Credit card merchants receive commission for each transaction, so increased transaction volume equals increased revenue. For brick & mortar banks, the major benefit is lowering the costs of branch offices where people usually do offline transactions.

[2] There is some discussion among FSPs of moving towards a model were end-users will become responsible for fraud, but considering Anderson (2001), such a form of liability dumping will most probably backfire, and the current strategy of the FSPs of keeping fraud at "acceptable" levels is best.

[3] Internet Corporation for Assigned Names and Numbers is one of the Internet's governing bodies: its responsibilities include allocating IP addresses and top level domain name management (including new country codes).

IP address 131.180.77.34. Knowing this IP address is necessary before the network connection can be established, and you computer obtains it by performing what is known as a *DNS lookup.* A method for blocking access to certain servers is by filtering traffic towards that particular IP.

Now, fast-flux abuses the DNS system to avoid being filtered, and also to perform some other neat devious tricks. This is how: as shown in Figure 20, in the fast-flux technique, bots continuously register and deregister their IP addresses for a particular domain. This means that each time a request is made by a user to access a malicious site, she will end up accessing a different bot (and possibly different ISP). This way, it won't be possible to take down the bot or blacklist its IP, as within a few minutes, a different bot will be serving the malicious content. The only option to combat fast-flux is for the registrar to suspend the bad domain.

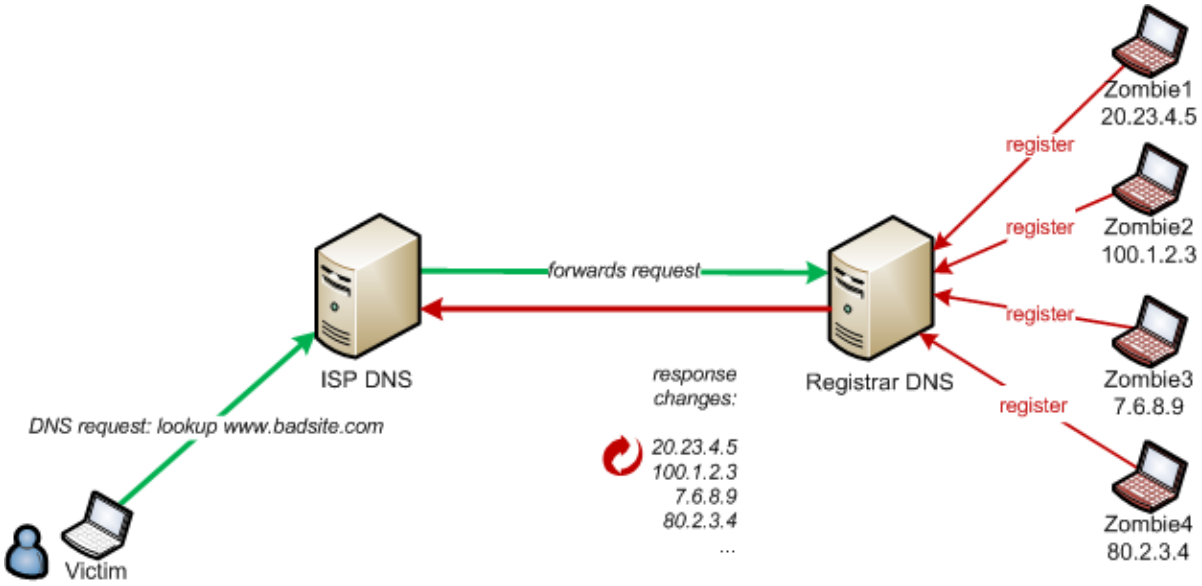

Figure 20 - Fast flux DNS technique, used by botnets to hide phishing and other malicious sites

In deciding whether or not to suspend a domain, registrars face several incentives. Positive incentives include: *maintaining reciprocity*, *avoiding brand damage*, and *avoiding being blacklisted*. Negative incentives include, most importantly, *legal constraints*. A case has to be built-up before a domain can be suspended (creating high costs); and the risk exists of mistakenly taking down a legitimate domain. An incentive that works both ways for registrars is *cost of abuse handling*. Domain registration is a very low cost, low margin business, making handling abuse complaints very expensive for the registrar. (Van Eeten and Bauer 2008)

These can lead to two extreme strategies. On the one hand, a registrar can choose to be proactive and even use automated abuse handling, to encourage criminals to move away from them and move to a more lax registrar; a totally opposite strategy is also possible: to not care AT ALL about abuse handling (targeting a different market segment in this case). In the majority of cases, registrars will adapt a strategy similar to that of ISPs: respond to external domain suspension requests (such as when submitted by a bank), and letting the problem exist if no one is complaining about it. Taking into account the number of phishing and malware domains not reported, there are externalities arising from the domain registrars' behaviour in the value-net. (Van Eeten and Bauer 2008)

INTERNET SERVICE PROVIDERS

ISPs choose among a relatively wide range of strategies when it comes to cyber-security. On one end we have ISPs that make money by doing absolutely nothing regarding mitigating cyber-threats, the so called *"rogue"* ISPs (Van

Eeten and Bauer 2008; Anderson et al. 2008a). Then we have various levels of vigilance, with most ISPs seemingly taking care of the most extreme cases of abuse on their networks, but not investigating every case. These strategies result from a mixed set of incentives, including some incentives that can work both ways (i.e., can have both a positive influence and negative influence on security).

Positive incentives

An ISP that has active spammers and spam bots running on its network, will surely end up receiving *abuse complaints* from other networks. ISPs need to act seriously on these abuse complaints, i.e., identify and stop the offending user and machine. This is time consuming and costly, but necessary to do, as a huge amount of security issues gets resolved via personal contacts and "informal networks of trusted security personnel". Thus maintaining reciprocity acts as an incentive for security. (Van Eeten and Bauer 2008)

ISPs that do not react in a timely fashion to their peers face the risk of getting blacklisted. A *blacklist* (also known as DNBL), is a method for tracking IP addresses of computers or networks linked to spamming. Most mail server software can be configured to reject or flag messages which have been sent from a site listed on a blacklist[1] (Wikipedia 2009a). If an ISP's mail servers get added to a blacklist, emails sent from those domains will be rejected. In other words, email sent by the users of that ISP will not get delivered, which can result in thousands of customer calls and complaints. This is a very costly punishment for an ISP, and hence an incentive for security. [2] (Van Eeten and Bauer 2008)

Negative incentives

A major disincentive for mitigating malicious users and infected users is *legislation ambiguity*. ISPs fear that by monitoring for malicious traffic, they will be breaking privacy laws. They also fear that they might be held liable for certain preventive measures they take, such as disconnecting infected users. (Van Eeten and Bauer 2008)

Difficulty in quantifying the costs/benefits of increased cyber-security is another disincentive. Management looks suspiciously at many security investments, not the least due to the overpromise of *"magic box"* security solutions in the past. (Van Eeten and Bauer 2008)

Incentives that are both positive and negative

An example of an incentive that can work both ways is cost of *customer care* and *call support*. As mentioned, not doing anything about security, and becoming blacklisted, could result into thousands of customer calls. Here, call support works as an incentive for ISPs to get serious on bots. On the other hand, getting too tough on infected customers also drives up the cost of call support (requiring communication with many customers). This time, the cost of call support acts as a disincentive for too much security. (Van Eeten and Bauer 2008)

Another example is the cost of *capital expenditure* (for infrastructure). Not taking security actions could translate into bandwidth being eaten up by spam, DDoS, and warez, which would necessitate investments in infrastructure expansion, and hence motivate ISPs do something about bots heads-on. On the other, in order to take action, it would be necessary to buy monitoring equipment and other security appliances, which can easily cost millions, and a reason not to do much . (Van Eeten and Bauer 2008)

Brand damage also work similarly in both ways: taking no action would create an image of untrustworthiness and bad citizenship, turning some customers away. Being too vigilant on the other hand, will also turn some away, as it

[1] There are quite a number of these block lists being run by volunteer organizations, such as SpamHaus, SORBS, and SpamCop.

[2] An average customer call or email can cost €8 - €16 , which destroys the profit margin of a user (Van Eeten and Bauer 2008)

creates fear that the ISP is too stringent and bothersome (i.e., puts constraints on users).[1] (Van Eeten and Bauer 2008)

Net effect

We have summarized all these incentive in Figure 21. It should not be surprising at this point why the majority of ISPs would fall somewhere in the middle, deciding to take action against the top 2% of botnet infections only, and acting in a reactive fashion mostly on incoming abuse reports by other ISPs. [2] Please note that the shape of the curve as given in this figure is only an example of how these incentives add up, and the exact empirical shape of the curve could be quite different.[3]

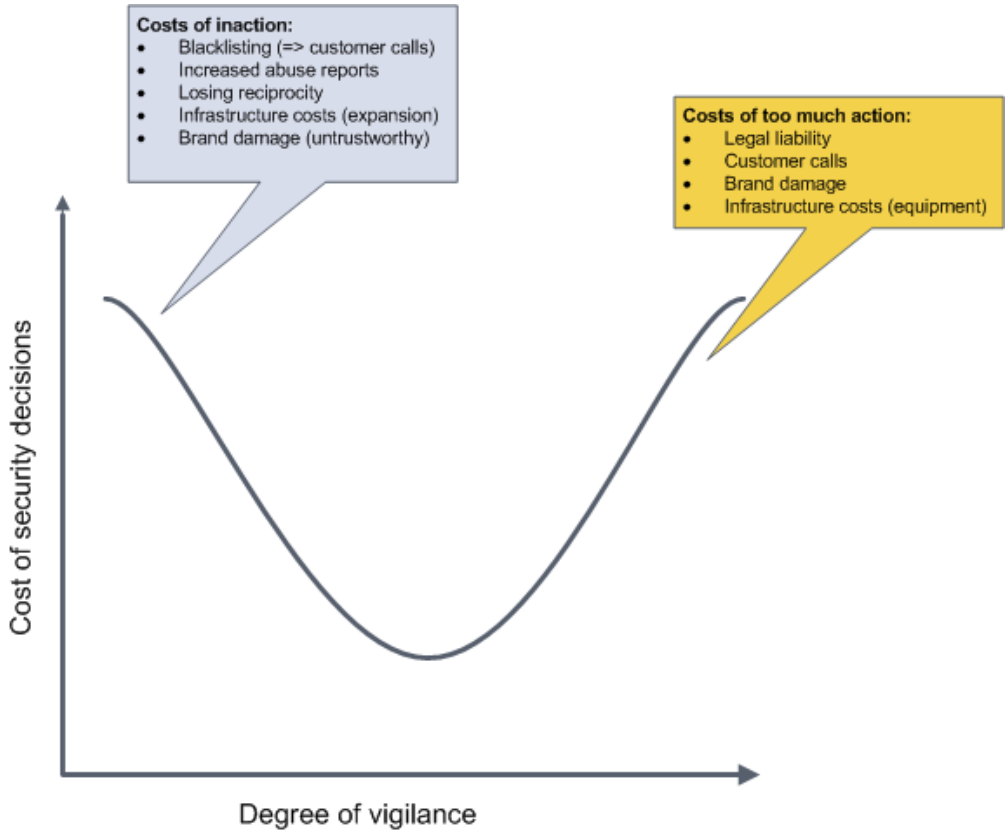

Figure 21 - Incentives that ISPs face in choosing their level of security (shape of curve is an example only)

How do these strategies add up in terms of externalities? It would not be accurate to state that ISPs would be creating externalities if they do not take action, as they are not the source (creators) of the network "pollution". But we can definitely state they are in a position to manage at least part of the externality being caused by end-users (and most of them are to some extent). There has been much debate and focus lately on ISP's responsibilities in this regards, which we shall visit in depth in an upcoming section. As a closing comment, it's good to know that despite the absence of any regulations, ISPs are doing increasingly more in terms of security.

[1] The strength of these incentives is not well known. For instance, it is not certain how strong a factor brand reputation is in customer purchases in the ISP market. Price competition seems to play a more important role many times.

[2] The rogue ISPs mentioned before decide to forgo support, abuse handling, reciprocity and branding all together, attracting a "particular" market segment.

[3] For instance, it could be closer to the classic exponential cost-benefit curves presented in economic literature.

Conclusion

In this section we looked at a number of the more important actors in the complex Internet ecosystem[1]. We present a summary of our findings in Figure 22, in regards to how each actor is contributing to the botnet problem (or alternatively, mitigating it); and how botnets are in return, affecting that actor.

Of course it should be noted that the real world is much complex then what we have presented in Figure 22. An important element missing in this figure is the inter-relation between the actors themselves, and feedback loops created in this regards. (For instance, the policies of an ISP would have impact on the behaviour of its end users, and vice versa, etc.). Considering these interactions in the actor analysis can be the topic of further research.

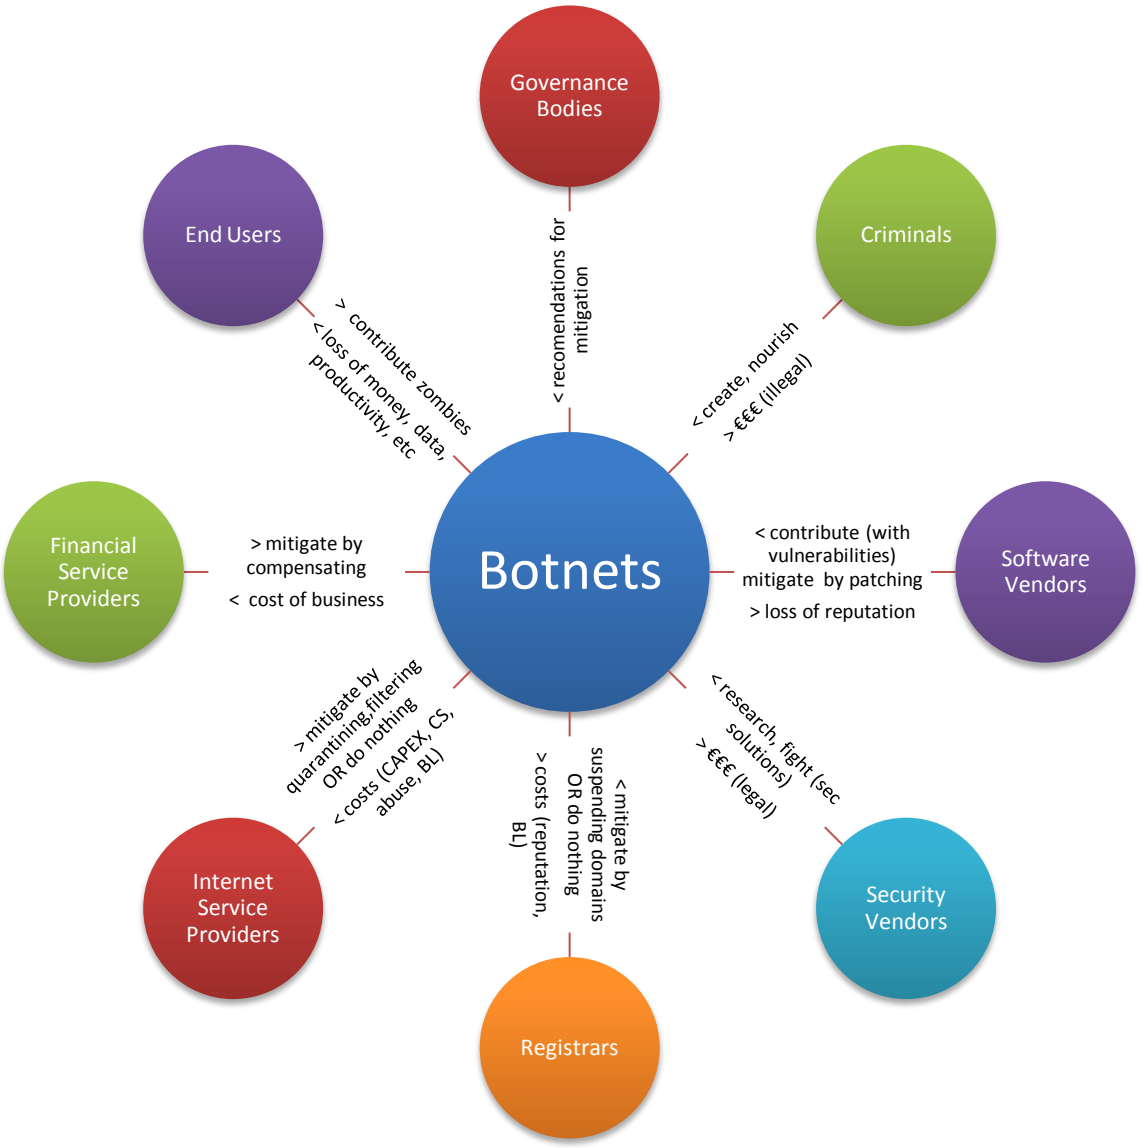

Figure 22 – Simplified diagram of relations of online actors to the botnet problem (without interactions)

[1] Some other actors could also be examined, such as Internet bodies (ICANN, IETF) , standard bodies, CSIRTS, ... These actors are influential in the cyber-security landscape, but to save space and as their roles are secondary, they have been omitted.

2.3.2 Mitigation via intermediaries

As we saw in the previous section, a wide variety of actors with different interests and strategies contribute to the botnet problem. Hence, botnets can be mitigated to some extent by focusing on each of the actors and its links with the problem (as shown in Figure 22). Some believe that among the different actors, incentivizing or pressuring Internet service providers to act is the most practical solution for botnet mitigation. In this section we shall review the arguments give by this group, and also by those opposing this view.

ISPs can mitigate the threat of botnets by adopting certain measures and procedures. For instance, they can use various techniques to detect bots within their network, notify their infected customers, and even help in 'remediation' (Livingood et al. 2009); they can stop harmful traffic going out of their networks (keeping the Internet clean), or coming into their networks (protecting their customers); and so on. What gives ISPs a unique position in fighting botnets, is, foremost, their role as provider of IP connectivity, which gives them the ability to act upon the bot traffic (see Figure 23), and secondly, their access to the end users (in the sense of contacting them, or asking them to install certain software, etc). Put in economic terms, ISPs have the capability to manage end-user *externalities*[1], should they have the incentives to do so.

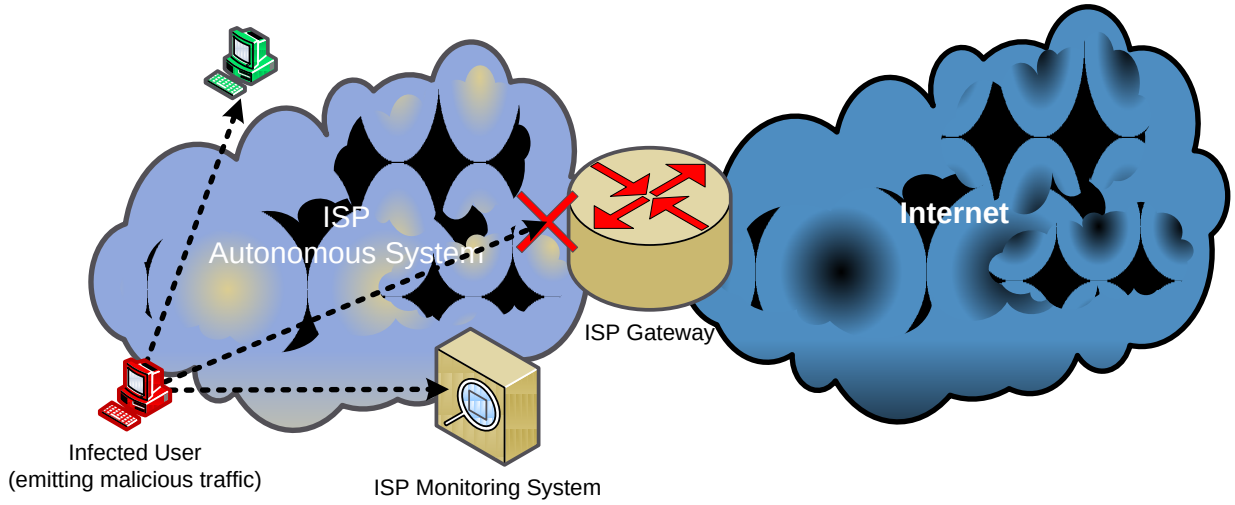

Figure 23 - Examples of how ISPs can detect and filter bot traffic

Anderson and his colleagues (2008b) are one of the proponents of making ISPs liable. In a report carried out for the European Network and Information Security Agency, they discuss various methods to fix externalities created by infected machines, and their final proposal is for the European Union to *"introduce a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines"* (see Table 9, recommendation 4) . Anderson and his colleagues ask the ISP industry to propose alternative practical methods for dealing with botnets, should they not agree with this proposal! [2]

Similar arguments are presented by others. In a 'request for comment' draft written for the Internet Engineering Task Force, Linvingood and others (2009) state that ISPs are in a 'unique' position for remediation of bots, and list reasons why taking steps in this regards would be beneficial for both themselves and their customers. They then

[1] In economics, an *externality* an economic transaction is an impact on a party that is not directly involved in the transaction

[2] They even state having a law enforcement officer stationed at each ISP as an extreme alternative, should ISPs be not willing to take on this responsibility themselves.

give recommendations on measures that need to be taken. Other industry groups, such as the Internet Industry Association of Australia (2009) and the Messaging Anti-Abuse Working Group (2009), have also released similar guidelines. In a recent EU conference on cyber-security, a much debated ideas was to make ISPs legally liable, for the damage caused by the data they transmit - at least to some extent (Economist 2009). Likewise, in this year's Virus Bulletin conference, a top Google executive called on ISPs to get tough on botnets:

> *"ISPs are the best place to do something about malware. They already have monitoring systems that could be used to identify signs of malware and botnet activity. If they see abnormally high e-mail activity, that's most likely spam from a botnet"* (Naraine 2009) .

Apart from ISPs having the unique ability to tackle botnets, focusing on them also makes the problem more manageable. Talking with a handful[1] of large ISPs will be more feasible for governments than holding millions of end-users responsible, many of whom even after cyber-security awareness and education campaigns will still lack the skills to ensure proper security (Wash 2007).

But is it fair to ask ISPs to take on this responsibility? Van Alystine believes that the current approach to filtering malicious traffic – which is to ask end-users to protect their systems with security software, is like letting factories freely pollute the environment, and then asking every man, woman, and child on Earth to wear a gas mask and boil their water before use [2] (GoogleTechTalks 2007). If we continue this analogy, we should be filtering network junk at the source, which would be sending ISPs. Of course this analogy is not totally correct, as first, ISPs could argue that they are not the creators of this junk; and second, in environmental pollution, the factories make money out of the pollution they produce, were as ISPs don't.

The legal grounds for holding ISPs liable are also somewhat shaky. In the U.K., in response to the governments suggestion of ISPs monitoring customers online, the Internet Service Providers Association (ISPA) stated that "the 2002 E-Commerce regulations defined net firms as *'mere conduits'* and not responsible for the contents of the traffic flowing across their networks" (BBC 2008a). There are also "transaction costs" involved, in case of lawsuits by customers disputing being cut-off from the net (Anderson et al. 2008b). Anderson and his colleagues make some discussion about this and add the clause "*...with a right for users to have disconnected machines reconnected by assuming liability"* to their original recommendation.

Others fear against ISPs taking on such a role, believing that some of the technology that would be implemented for identifying bots, most notably *'deep packet inspection'*, could end up being used for hostile purposes, such as data mining, eavesdropping, censorship, and voiding 'network neutrality' (ArsTechnica 2007; NetEqualizer 2009) . Regulation and oversight might be needed to avoid such applications.

The reality is that already many ISPs are taking substantial steps in regards to malware and bot infections. Some have implemented technologies that automate the process of monitoring malicious behaviour on their network and quarantining infected machines (Van Eeten and Bauer 2008). The real issue is the *extent* to which this is done. As discussed in the previous section and depicted in Figure 21, ISPs have incentives to not totally let go of the malware issue, and yet not to try to completely fix the issue. In economic jargon, ISPs will tackle botnets up to a point where the 'marginal cost' of being more secure exceeds the 'marginal benefit' of being more secure. This could be much lower than the social optimum. Van Eeten and Bauer (2008) quote a security expert saying "unless ISPs are contacting more than 10% of their customer base on a monthly basis, they are effectively taking no

[1] Even in larger markets such as the United States, the number of "true" ISPs (not virtual or resellers) is in the order of tens of companies (source: TeleGeography database)

[2] More precisely, Van Alystine comment is about spam emails, but the idea is the same

action". But given the *scale* of the potential risk of botnets, it is hard to see how any ISP could cope with complete liability (Economist 2009).

A RESEARCH GAP

Given all the arguments, focusing on ISPs remains a viable and practical solution for fixing the botnet problem; however, doing so has certain hurdles that need to be overcome. A question that comes to mind is whether regulation is needed to convince ISPs to tackle the problem, or would it be better to pursue incentives that would make the market self-regulate? To answer these questions we would need to look at the incentives currently influencing ISPs.

As previously mentioned, Van Eeten and Bauer (2008) have performed a qualitative study elaborating on the incentive structure of online market players, when it comes to security decisions. A mixed incentive structure was identified for ISPs. A *quantitative* study in this area could prove fruitful to better understand the strengths of the incentives and how they combine with each other, answering questions such as:

> *... Do ISPs significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? Can we identify internal or external factors that can explain this variance? ...*

In their recommendations to ENISA, Anderson and his colleagues (2008b) correspondingly argue that a prerequisite for designing and adopting policy actions to align incentives for improving information security is to have quantitative data (with consistent metrics) regarding the security performances of actors. Based on this argument, they recommend that data about the quantity of spam and other bad traffic emitted by ISPs be collected and published (see recommendations 1 to 3 in Table 9). They state that "ISPs (and banks) are two particularly problematic 'black hole' where data are fragmentary or simply unavailable".[1]

A unique opportunity exists for undertaking such a research in the TPM faculty of TU Delft. The faculty has access to a database consisting of billions of spam email messages, sent from 30,000+ autonomous systems, over the last few years (2005-2008). This unique dataset can be used as the basis of a quantitative study into the role and incentives of ISPs in mitigating botnets. This is because spam can be used as a proxy for botnet activity. (This idea will be elaborated on, in the methodology chapter.)

Using this quantitative data and statistical methods, it would be possible to uncover the extent that certain incentives and factors influence what ISPs do (in regards to security), and how those actions in turn, influence botnets. It would also be possible to identify best and worst performers, and take a look at what separates them from the pack.

Such a quantitative study into the role and incentives of ISPs regarding botnet mitigation, will fill in a gap in the scientific literature, and also aid in policy making.

[1] From an economic perspective, this is a problem of 'incomplete information' for the policy maker.

2.3.3 LIST OF SECURITY MEASURES

Up to now we have talked very generically about the various security measures that Internet service providers can adopt. But how broad is this range of security measures?

We discussed that each ISP is an autonomous system that chooses its own set of rules, not one that is dictated by a central authority, and many times there is no consensus on security measures that work best. We also discussed that the ISPs have a mixed incentive structure, and varying levels of security performance (section 2.3.1). These facts together imply that ISPs have a large number of choices available when designing their security policies and adopting their security measures. In this section, we will attempt to make a fairly exhaustive list[1] of such security measures.

We have extracted measures from several industry driven efforts to come up with a set of best security practices, and one survey of the measures that are really in use. The sources are as follows:

- Provider Security Measures, *survey* (ENISA 2007)
- Best Practices in Anti-Spam (ETIS 2007; MAAWG 2005; Schryen 2007 ch. 4; Sendmail 2007; OECD 2005)
- Best Practices in Anti-Phishing (MAAWG and APWG 2006)
- Best Practices for Mitigating Bot Infections (MAAWG 2009, 2007a; Livingood et al. 2009)[2][3]
- General Best Practices for ISPs and Network Operators (MAAWG 2007b; IndustryCanada 2005)

The number of measures listed in these sources goes to over 200, and they include both *technical* and *organizational* aspects. We have grouped them into nine major groups, and listed them in Table 10. The major groups are as follows:

- Active abuse handling
- Proactive detection of malicious activity
- Filtering malicious traffic and content
- User education and awareness
- Client security and quarantining
- Using updated network protocols and servers
- Participation in the security community
- Management and administrative procedures
- Legal measures

Please note that the explanation of the measures is out of the scope of our work. The point here is to show that the proposed and adopted list of security measures is very large. Interestingly, the effectiveness of many of these measures is not solidly known. They are all debatable, depending on context and the extent to which they are implemented. The specific choice and extent of measures will in the end be, without doubt, influenced by the particular incentives of each ISP.

[1] Please note that our list of measures is regarding the security policies that effect malware, spam, and botnets. Security has many more aspects as well.

[2] MAAWG is the *Messaging Anti-Abuse Working Group.*

[3] The work by Livingood and others is a memo is published by the *Internet Engineering Task Force*

Table 10 - List of security measures proposed to ISPs in the literature

Measure Category	Specific Measures
Active abuse handling	• Provide contact details for email abuse and security violations • Monitor RFC2412 addresses (abuse@domain, etc) • React to complaints from other ISPs about security and spam (and track them) • Respond to subscriber complaints about spam • Abuse desk automation (using in-house system, ARF, and feedback loops) • Keep public records of all publicly routable/visible IP addresses, and domain names (such as WHOIS, reverse DNS, SWIP, etc) correct, complete, and current.
Proactive detection of malicious activity	**Bot detection:** • Monitor traffic peaks • Monitor email bounces • Actively monitor the volume of inbound and outbound email traffic to determine unusual network activity and source of it • Botnet detection via DNS • Botnet detection via IP space scanning • Botnet detection using NetFlow **New threat detection:** • Deploy real-time traffic anomaly and/or signature based detection mechanism • Use of blackholing and sinkholing to secure services • Deploy spam-traps to optimize anti-spam installations • Use in-house or 3rd party 'security intelligence service' • Analyze where spam comes from • Communicate and share data via feedback loops ?
Filtering malicious traffic and content	**Basic filtering (ingress and egress):** • Block inbound port 25 (spam-relay) • Block inbound port 53 to residential customers (avoid fast-flux) • Manage access to outbound port 25 for hosts on residential network (spambots) • Drop egress spoofed IP sources **Content filtering (ingress and egress):** • HTTP: provide proxy service to filter bad web domains (phishing sites, etc) • SMTP: anti-virus scan and spam-filters on network (e.g. during DATA phase) • Block potentially infecting email attachments • Inbound filtering of phishing messages • Email content scanning: Bayesian filters, heuristic, probabilistic, frequency analysis, fingerprinting, URL-based, etc • Disconnecting SMTP connections w/ unknown recipients • Change incoming mail: disable hyperlinks, hide images from untrusted sources ? • Outbound content filters (perform virus-scanning for outbound email,…) **Dynamic filtering systems (based on IP reputation):** • Dynamic IP based sender reputation (also known as real-time-blacklists) - at IP level, or SMTP level • Add offending subscribers to blacklist • White-listing (ISP mail-servers, good customers) • Short-lived blocks of web-traffic for suspicious sites **Slowing suspicious traffic:** • Use of traffic shaping as a security method

Measure Category	Specific Measures
	• Grey-listing • Slowing down SMTP connection • Limit the volume of outbound mail • Set greet pause in MTA *Note: 'Ingress' traffic is applied to traffic coming from outside, where as 'egress' filtering is applied to traffic leaving the network.*
User education and awareness	**Education, training, and campaigns:** • Provide information on security via website or email (and other channels) • Inform subscribers of risks of not implementing counter measures • Provide educational literature to users with best practices for avoiding malware • Use of customer portal for information • Communicate security policies and procedures to subscribers • Provide to users links for educational resources regarding nature and scope of threats **Help in remediation:** • Notify users via email/telephone/walled-garden/in-browser/IM/SMS/ of infection • Maintain a well-publicized security portal where a compromised user can be directed for remediation • Provide security tools, education and useful links for users to perform own remediation • Inform subscribers of costs of remedies and also point to professional help • Detailed guidance to subscribers (provide a guided flow for remediation process) • maintain forum for users (self-help) **Customer Call Support:** • Train call centre agents on how to assist users • Abuse department customer dialog to help in disinfections
Client security and quarantining	**Disconnecting and quarantining infections:** • Place infected users walled gardens (based on abuse report, or internal detection) • Quarantine computer in networks unless protected • Allow escape from walled garden if trusted, or if certain software is installed • Service suspension on repeated security failures • Service termination for non-compliant subscribers *Walled garden measures* – access lists; redirect HTTP; redirect botnet C&C to honeypot; manage outbound SMTP to quarantine/honeypot; **Securing end user Clients:** • Provide security software (spam-filtering, anti-virus, browser plugins, etc) for clients, and encouraging their use. This can be for free or for a reasonable price • Provide subscribers information about availability and use of such solutions with links • Provide NAT routers with firewalls to customers • Subject users to mandatory scan (when first provisioned, or periodic)
Using updated network protocols and servers	**Simple Mail Transfer Protocol:** • SMTP Authentication: SMTP AUTH, TLS, Pop3 before SMTP,… • Provide message submission for mail and ensure only account holders use it • Use FQDN in EHLO/HELO • Use sender validation (DKIM/Sender-id/SPF/reverseMX) on inbound email • Reject email if detected forgery w/ sender authentication

Measure Category	Specific Measures
	• Prohibit sending of email with forged headers • Configure human readable delivery status notifications (?) **Other:** • Implement DNSSEC • Ensure DNS architecture is up-to-date (to avoid cache poisoning)
Participation in the security community	**Membership:** • Become member of an industry association • Join one or more anti abuse forums • Compare effectiveness of anti-spam installations with others? • Share information and data on the intensity and scope of spam and its evolution • Methods for sharing dynamic IP address space information with others **Notification:** • Communicate knowledge of phishing attacks to the targeted institution • Contact an ISP directly when receiving spam from it • (allow spam source time to solve the problem before blocking traffic) • Share evidence of bot with remote sites • Implement / use feedback loops (e.g., between ETIS partners)
Management and administrative procedures	**Ensuring security level by adhering to:** • Industry best practices • National legislation guidance • International standards (ISO 27002:2005, ISO 27006:2007) • Using SLAs to ensure appropriate level of security **Formal panning:** • Business contingency plan for protection of network integrity • Disaster recovery plan for protection of network integrity • Annual testing of business continuity plans • Use a risk management process **Other:** • Constantly improve knowledge and operating practices • Review anti-spam installations for common practices? • Multilevel abuse handling ? • Build necessary tools for care agents to retrieve relevant info (about bot detection) • Protect customer email addresses? • Written security guidance for staff and subscribers • Train support representatives about fishing and scams
Legal measures	• Adopt and enforce Acceptable Use Policy (AUPs) • Forbidding spamming in Terms and Conditions • Informing subscribers of legal consequences of sending spam • Inform NRA of security breach • Inform customers / public of security breach • Report spam to NRA (national authorities) • Pursue legal actions for spam

2.4 Building a Conceptual Framework

In the final section of this chapter, we would like to set the setting for the rest of the thesis. We will start by summarizing briefly what we have learnt so far.

2.4.1 Formulating the research question

We started by enumerating the online security threats that end-users face, and saw that many of those threats have a wider circle of influence than just the victim. Rather, victims are used as a stepping stone for the next phases of complex attacks. This is particularly true about botnets. We then shifted solely onto botnets – how they're formed, commanded, and used. We recognized a sense of urgency to tackle botnets, and by reviewing the literature on the actors involved in the arena, focused on ISPs as an intermediary with certain powers in mitigating bot activity and propagation.

We then looked at the heated discussion surrounding the role of ISPs in the "war" against botnets, and identified a research gap in the literature regarding this discussion: the extent to which ISPs are already mitigating botnets is not fully known. And although various incentives that shape an ISP's security decisions have been identified, the strength of these incentives in adopting particular measures is not understood; neither is the effectiveness of the adopted security measures. Knowing the answers to these questions is crucial for the debate. They will serve to deepen the scientific understanding of the botnet phenomenon, and also aid policy makers in choosing the correct combination of 'carrots and sticks' to remediate the situation.

This all leads us to the following problem statement (partially mentioned):

> ***Problem statement:*** *Are Internet Service Providers crucial intermediaries in botnet mitigation efforts? Do they significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? And what implications does this have for policy?*

To answer this research question, we would need to answer a set of sub questions:

> ***SubQ1:*** *What are botnets, and why is it important to mitigate them?*
>
> - *What are the major security threats the Internet faces?*
> - *What makes botnets standout as the most serious of threats?*

> ***SubQ2:*** *Who are the main actors that can mitigate botnets? Are ISPs the key intermediary for such efforts?*
>
> - *Who are the major actors involved in botnets?*
> - *What are their incentives to adopt specific strategies? What externalities do they create and absorb?*
> - *Among these, what makes ISPs an interesting candidate for mitigation efforts?*

SubQ3: *Do ISPs significantly differ in the degree in which they mitigate botnets?*

- *How can we quantifiably measure ISP security effectiveness (in mitigating botnets)?*
- *How different are ISPs (in terms of magnitude of botnet activity on their networks)?*
- *What is the list of security measures that ISPs choose among? How wide is their choice?*

SubQ4: *To what extent can we explain the varying degree in which ISPs mitigate botnet activity? Can we identify internal or external factors that can explain this variance? (This question has both quantitative and qualitative aspects)*

- *Example external factors includes end-users, national infrastructure, and criminals*
- *Example internal factors includes business strategies, adopted security measures, etc*
- *In place of factors, we might be able to identify certain characteristics of ISPs that explain the variance*
- *Can incentives explain what we are identifying?*

SubQ5: *What are the implications of the above findings in terms of practical policy options for botnet mitigation?*

Please note that sub-question 1 and parts of sub-question 2 have already been answered as part of the literature review; the rest of the sub questions will be answered in the upcoming chapters.

2.4.2 The conceptual framework

During this chapter we saw that since cyber-security is a relatively new problem, there isn't always a clear consensus or understanding of the factors at play and the causal relations. In the current literature, there isn't what we could call "firm" hypothesis, i.e., clearly stated hypothesis that could be simply copied into a conceptual framework (except maybe for the fact that economic incentives need to be addressed, which most authors agree.)

Nonetheless, it is possible to generate an initial sketch of the relations that hold and what influences what, based on the different sources. In Figure 24, we present a conceptual framework of factors, actors, and incentives that effect botnet activity at the ISP level (which is our unit of analysis). We must iterate that this is only an initial attempt at making hypotheses regarding this topic. Please note that not all of the relations will be easy to investigate, so in the methodology chapter, we shall come up with a set of 'empirical' hypotheses based on operational feasibility. The provided framework is further explained below

Major relations

Table 11 lists the major relations present in the framework. The starting point of the framework is that botnets are formed and commanded by criminals; hence we can recognize them as the major cause (R2). But criminals rely on the risky behaviour of end-users to perform their activities, so the behaviour and type of an ISP's customers (on average) will greatly influence botnet activity as well (R1). The existence of unpatched software and increases in broadband connections, help the bots in propagating and activity (R3). However, as discussed in section 2.3.2 and 2.3.3, ISPs can undertake various security measures that will mitigate the effects of botnets (R4). The security measures that an ISP adopts, is related to its mix of incentives and the costs/benefits it perceives (R5). This last relation is somewhat weakened by ambiguity in what is legal and effective (R6). Among the relations, R1 and R2 are direct causal factors (as causes of botnet activity), where as R3 and R4 are moderating factors.

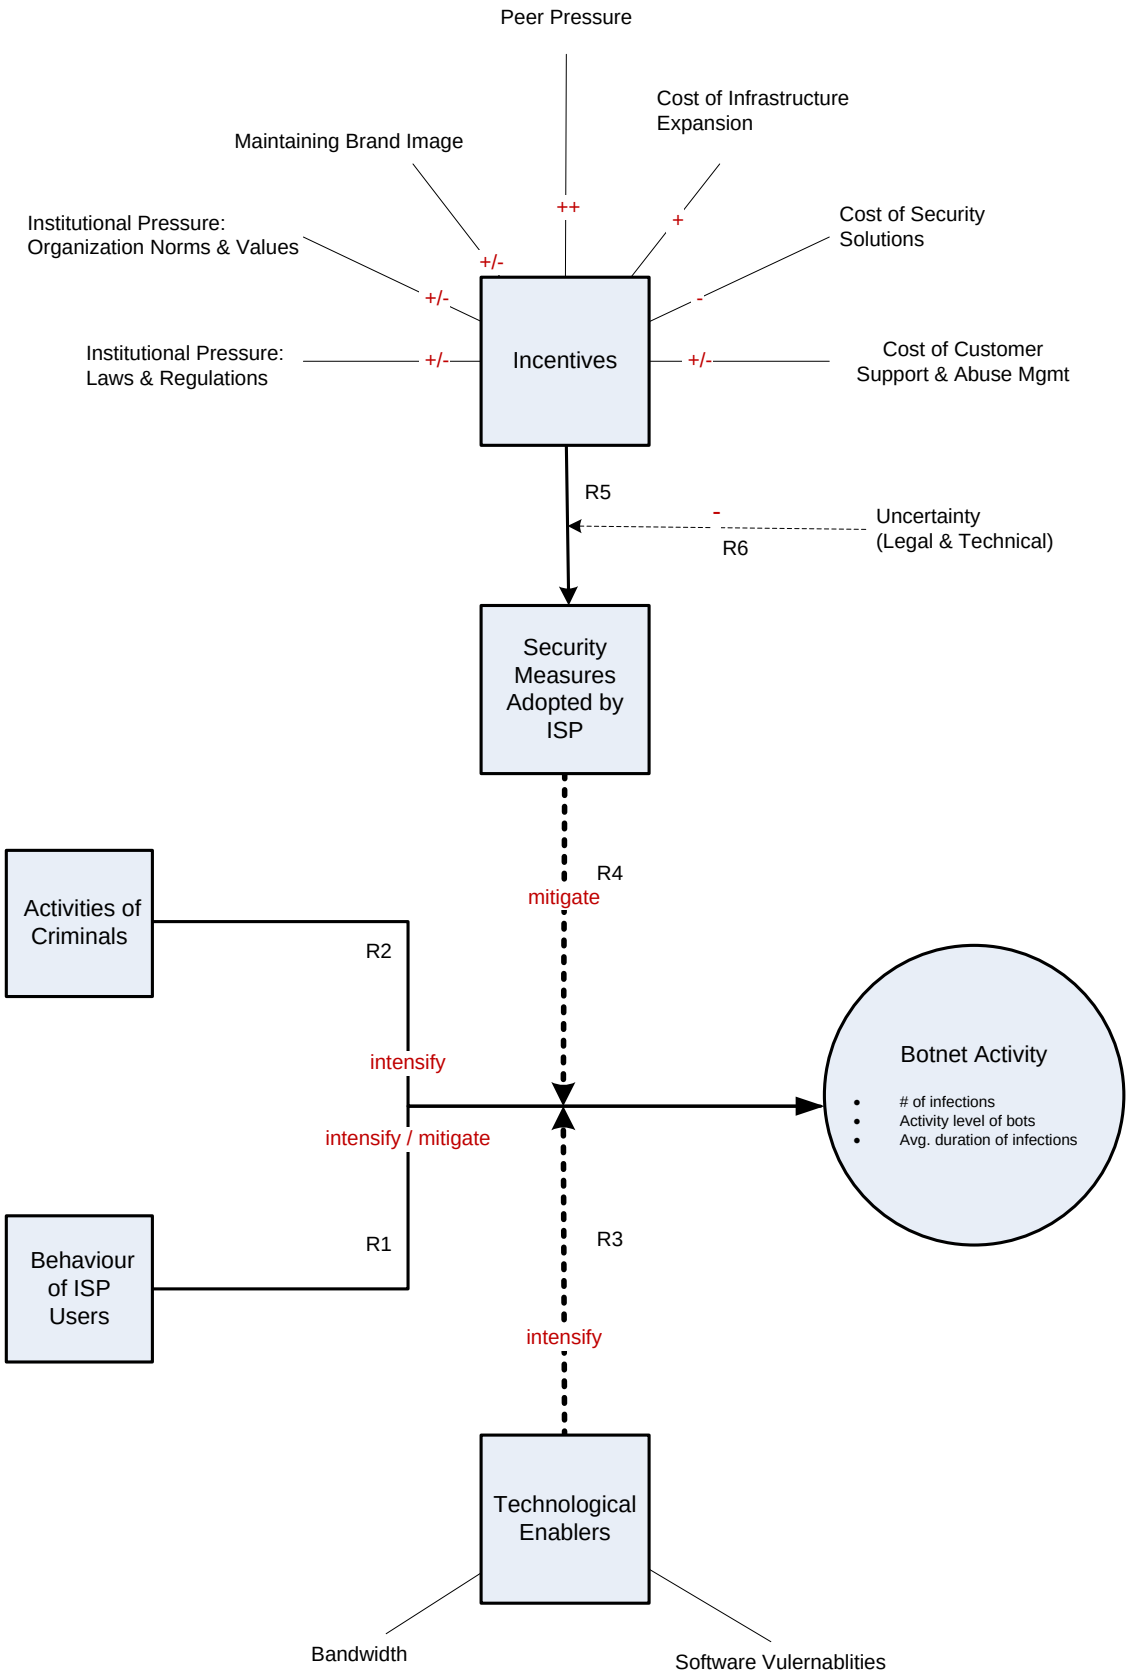

Figure 24 - The conceptual framework of factors influencing botnet activity at the ISP level

Table 11 - List of factors that influence level of botnet activity in an ISP (see conceptual framework)

Factors Influencing Botnets Activity	Examples/Notes	Effects on Botnets	Source
R1- Behavior of ISP users (on average)	• Education & Awareness • Attitudes • Available Technology (UI design, etc) Influenced by market segment and demographics	*Intensify / Mitigate*	(Van Eeten and Bauer 2008) (Wash 2007) (OECD 2007) Others
R2- Criminal Behavior	• Based on Intent, Skills • Very dynamic Note: Opening this box is out of our scope.	*Intensify*	(OECD 2007) (TeamCymru 2006b) (Bauer, Van Eeten, and Chattopadhyay 2008) Multiple industry reports
R3- Technological Enablers	• Bandwidth (on average) – e.g., broadband speed • Software vulnerabilities (time period, services)	*Intensify*	(OECD 2007) (Anderson et al. 2008b)
R4- ISP Security Measures	• Technical & Organizational • Nine major categories (presented in section 2.3.3) Please note that effectiveness of individual measures is also out of our scope.	*Mitigate*	(Van Eeten and Bauer 2008) (Anderson et al. 2008b) (Economist 2009) Multiple best practices (in 2.3.3) Others
R5- Effects of Incentives on Security Measures	Causes different levels of security and choices of measures to be adopted	*Indirect*	(Van Eeten and Bauer 2008) (Anderson 2001) (Anderson and Moore 2007) (Anderson et al. 2008b) Others
R6- Uncertainty	• Legal ambiguity • Technical uncertainty: ambiguous (cost) effectiveness of various security measures	*Indirect*	(Van Eeten and Bauer 2008) Others

ISP incentives

Table 12 lists the main incentives that influence ISP's in adopting their certain level of security, i.e., the trade-off between security and other factors that they find most economical. As explained in previous sections, most of these incentives have been identified through interviews with stakeholders, but we have added incentives from the '*New Institutional Economics'* literature. In this literature, *norms and values*, and *laws and regulations*, are mentioned as two institutions that influence the rules of the game and the decisions players make (Koppenjan and Groenewegen 2005). Hence they have been added to the framework, although in particular, 'organizational culture and values' was not directly mentioned in the interviews[1].

As you recall from section 2.3.1 (in Figure 21), most ISPs end up adopting a security level in the middle – not too much, not too little. This is reflected in the table - the 'mixed' effects of incentives on security. The strengths of many of these relations are not well known, and would be interesting to discover.

Final words

In this chapter we reviewed the literature on cyber-security, and funnelled our way towards botnets, and their mitigation via Internet service providers. The synthesis of the literature was presented in this final section, in the form of research questions and a conceptual framework. These will be used to design our research in the next chapter.

[1] Consider an ISP where a famous security "hotshot" works: security is firmly valued by the technical staff, and they have considerable influence over the business units; hence the level of security goes up. In a "rogue" ISP, the reverse is true.

Table 12 - List of incentives that influence security decisions made by ISPs (in regards to botnets - see conceptual framework)

Incentive Influencing Security Decisions	Examples/Notes	Effects on Security	Source

Institutional Pressure: Laws & Regulations	• Regulations to abide to
	• Liabilities to avoid, e.g.:
	  ○ Invasion of privacy
	  ○ Damage due to attacks
		Mixed	(Van Eeten and Bauer 2008)
			(Anderson et al. 2008b)
			(Koppenjan and Groenewegen 2005)
			Others

Institutional Pressure: Organization Norms & Values		*Mixed*	(Koppenjan and Groenewegen 2005)
Maintaining Brand Image	• Costs of reputations effects and brand damage • Costs of customer acquisition Influenced by customer expectations (market segment)	*Mixed*	(Van Eeten and Bauer 2008) Others
Peer Pressure	• Cost of blacklisting (indirect cost) • Benefits of peering agreements • Benefits of maintain reciprocity	*Positive (strong)*	(Van Eeten and Bauer 2008) (Anderson et al. 2008b)
Cost of Infrastructure Expansion	Network expansions costs (also current spare capacity)	*Positive (weak)*	(Van Eeten and Bauer 2008)
Cost of Security Solutions	• Cost of implementing security solutions • Cost of capital (or availability of funding)	*Negative*	(Van Eeten and Bauer 2008) Others
Cost of Customer Support & Abuse Mgmt	• Costs of customer support • Cost of abuse management Influenced by customer expectations	*Mixed*	(Van Eeten and Bauer 2008) Others

CHAPTER 3- RESEARCH METHODOLOGY

3.1 MEASURING SECURITY EFFECTIVENESS

Introduction

The goal of this chapter is to present the methods that we will employ to answer our problem statement. The problem statement that we presented in the previous chapter – after narrowing down our topic and identifying a research gap, was as follows:

> ***Problem statement:*** *Are Internet Service Providers crucial intermediaries in botnet mitigation efforts? Do they significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? And what implications does this have for policy?*

We stated that we wish to answer this question quantitatively, and in this chapter we will discuss how we intend to do this. The key variable to measure is ISP security effectiveness (in terms of bot mitigation); Later, the variance of this variable among different ISPs can be analyzed, and seen how much of it can be explained statistically with the various factors external and internal to the ISP.

We shall operationalize the measurement of ISP security effectiveness by using *outbound spam* as a proxy. This will be done so by processing a large dataset of spam emails. Let us start this chapter by explaining the origin of this data, and also answer the question of why outbound spam is a valid proxy. Later on the dependent and independent variables will be discussed, and finally, a set of empirical hypotheses will be presented.

3.1.1 DAVE RAND'S SPAM TRAP

Dave Rand is a word leading expert on Internet security, and a renowned fighter of spam. He is the co-founder of MAPS, the first anti-spam blacklist on the Internet, with its roots going back all the way to 1996 (Wikipedia 2009d).[1] For a decade or so, Rand has been operating what's known as a *'spam trap'*, tracking the behaviour of spammers, and logging in real-time (a subset of) the spam activities on the Internet. The spam logs of this trap today constitutes of Terra Bytes of data and billions of spam messages! Rand has kindly provided a condensed version of this data to the author's research group in the TPM faculty.[2] This version contains logs of the spam sending incidents, but not the actual contents of the spam messages (more information shortly). The data, which is a time series from 2005 to 2008, consists of nearly 1 billion records and is approximately 100 GB in size.

HOW A SPAM TRAP WORKS

A spam trap is a mail-server that's only purpose is to receive spam (and catch all forms of it, from commercial to phishing and malware). The idea is to have mail domains and email addresses that do are not intended to receive legitimate emails. These addresses are then posted on websites and other places that spammers typically *"harvest"* (collect) email addresses from - to add to their gigantic contact lists. (Recall that the definition of spam is

[1] The parent company of MAPS was acquired in 2005 by the security firm Trend Micro (MAPS 2005). Since that time, Rand has worked in Trend Micro, and currently holds the title "Chief Technologist of Internet Content Security".

[2] Governance of Infrastructures, in the section of Policy, Organization, Law, and Gaming

unsolicited bulk emails). Spam traps log the actual spam message, the time it is received, the IP addresses of the spam sending machine, and other details of the connection (Shadowserver 2007b).[1]

OUTBOUND SPAM AS A PROXY FOR BOTNET ACTIVITY

In the late 1990s and early 2000s, spam was sent through the infrastructures setup and owned by spammers, and through their own Internet connections. This was partly due to the fact that spam wasn't totally illegal at the time (think of it as a grey business). Early this millennium, when strict laws regarding UCE came into effect, and at the same time, anti-spam measures based on '*IP reputation'* became mainstream, the spam operators faced serious trouble. Owning your own spam-servers could get you into legal trouble (and there were several high profile cases). Even if the servers were located outside the jurisdiction of the laws, the range of IP addresses belonging to the spam-servers could get blacklisted, meaning that no provider or business would accept emails from them.

Spammers, not to be so easily out done, turned towards the devious idea of sending spam via other people's machines – hence, the rise of the botnets. (This is actually a very interesting example of both legal and technical measures failing, due to the fact that the underlying economic incentives were not addressed.) This trend caught-on in late 2004. An era started in which the majority of spam is sent via *spam-bots*.

Estimates vary, but experts believe that 75 to 90 percent of spam today comes from botnets. (The rest, such has "snow shoe" spam, is sent via other methods; for instance, from rogue ISPs and using "hijacked" IP address ranges.) [***] The percentage differs based on the domain of the recipients. Our spam-trap belongs to an end of spectrum where the majority of the spam received has been sent via botnets. This is because this particular spam trap is a small and old domain. Spammers who are still sending spam form their own servers are much more careful in how they allocate their resources (as their costs are higher), and usually go for more targeted campaigns and fresher contact lists. (The email lists used on the bots on the other hand not sensitive and contain more stale records). Hence, we can safely conclude that the spam activity recorded in our spam logs are a good proxy for botnet activity.

There are still three points to consider: first, not all bots are busy sending out spam (recall that in the case of the Storm botnet only a fraction were actively sending spam at a time); second, our logs capture only a part of the global spam-bot activity, only a portion of the spam would be addressed to our spam trap. These are however not major problems, as they can be merely issues of sample size. As long as our data is consistent with the global spam/botnet trends, and only smaller in size, then statistical analysis will bear the correct results. This issue will be tested in Chapter 4, by triangulating our data with other publicly available sources.

A third and more serious error originates from certain technical limitations, such as dynamic IPs, NATs, and port 25 blocking. We will shortly discuss how these issues effect us, but before we do, the process of creating the dependent variables needs to be discussed.

In short, we believe that Dave Rand's database can be used to validly and reliably measure ISP security effectiveness (regarding botnet activities). It will however not be perfect, and we ne need to accept this as a limitation of our research, and take care in generalizing our final conclusions; a point which we will reiterate multiple times during this chapter.[2]

---

[1] Spam traps have a close cousin called *'honey pots'* which use similar techniques, but with the aim of gathering information on cyber attacks, worm outbreaks, etc.

[2] In a workshop held in September 2009 with Internet security experts from the Dutch ISP XS4ALL and the anti-spam firm Cloudmark, the experts approved of our measurement tool for botnet activity.

3.2 BUILDING THE DATASET

In this section we will explain how we plan to build the dataset that will be used for this research. This section will be mostly conceptual; some of the concrete steps will be discussed in the next chapter (data preparation).

3.2.1 THE DEPENDENT VARIABLE(S)

As outlined in the conceptual framework (presented in the previous chapter), our dependent variable is some indicator of botnet activity on an ISP's network. We propose the following three metrics for measuring botnet activity:

- The number of infected bots in the network in a specific period of time
- The amount of malicious traffic emitted by those bots during the period
- The average duration of days a machine remains infected

These metrics are all measures of different aspects of botnet activities. Let's start with the first metric: all else being equal (*ceteris paribus*), we would expect a more 'vigilant' ISP (that is, a more caring and effective one in regards to security) to have fewer bot infections on their network. As an example, consider two almost identical ISPs (in terms of operating conditions, user-base, environment, security measures, etc) - except that one of them provides free anti-virus software and security brochures explaining online risks to its customers. We would expect to have a lower number of bot infections in the ISP that is taking these extra steps. [1]

The second metric points towards measures an ISP takes that would limit the amount of damage a bot inflicts on the rest of the world, *post*-infection. Again, consider the two identical ISPs, with one of them implementing outbound content filtering. We would see lower malicious traffic with the same number of infections.[2] Tackling only the worst offenders (instead of implementing general security procedures) will also lower this metric.

The third metric is influenced by thing such as if an ISP notifies infected end users and helps them in the remediation process. A variation of this third metric is the average infection duration of the *worst offenders* (i.e., bots emitting very high amounts of spam).The rationale behind this variation is that the minimum an ISP should be doing regarding security is taking out the worst offenders. As you recall from chapter 2, ISPs do not go after resolving all cases of network abuse, but rather target the ones that others complain about, which would typically include the most spam-emitting bots.

Based on all these metrics, it would be possible to rank ISPs on their security merits. A word of caution must be raised on interpreting the first two metrics: they are absolute numbers and do not take into account the size of the ISPs (we will revisit this concern shortly). Interestingly enough, Anderson and his colleagues (2008b) suggest the collection of security metrics for ISPs that are much similar to the ones we have just proposed.

BUILDING THE DEPENDENT VARIABLES

The format of the spam logs available to us is presented in Figure 25. There is one log file for each day (adding up to over 1400 log files). Each of the files contains the following values of interest for us:

[1] Unless we have some reason to believe that the bad guys are targeting one of the ISPs in specific.

[2] Note that blocking outgoing malicious traffic is controversial among security experts, as some say that it is like sweeping the dust under the carpet, without remediating the infected bots

- The IP address of the spam sender
- The count of spam messages sent by that IP during that day
- The Autonomous System Number [1]

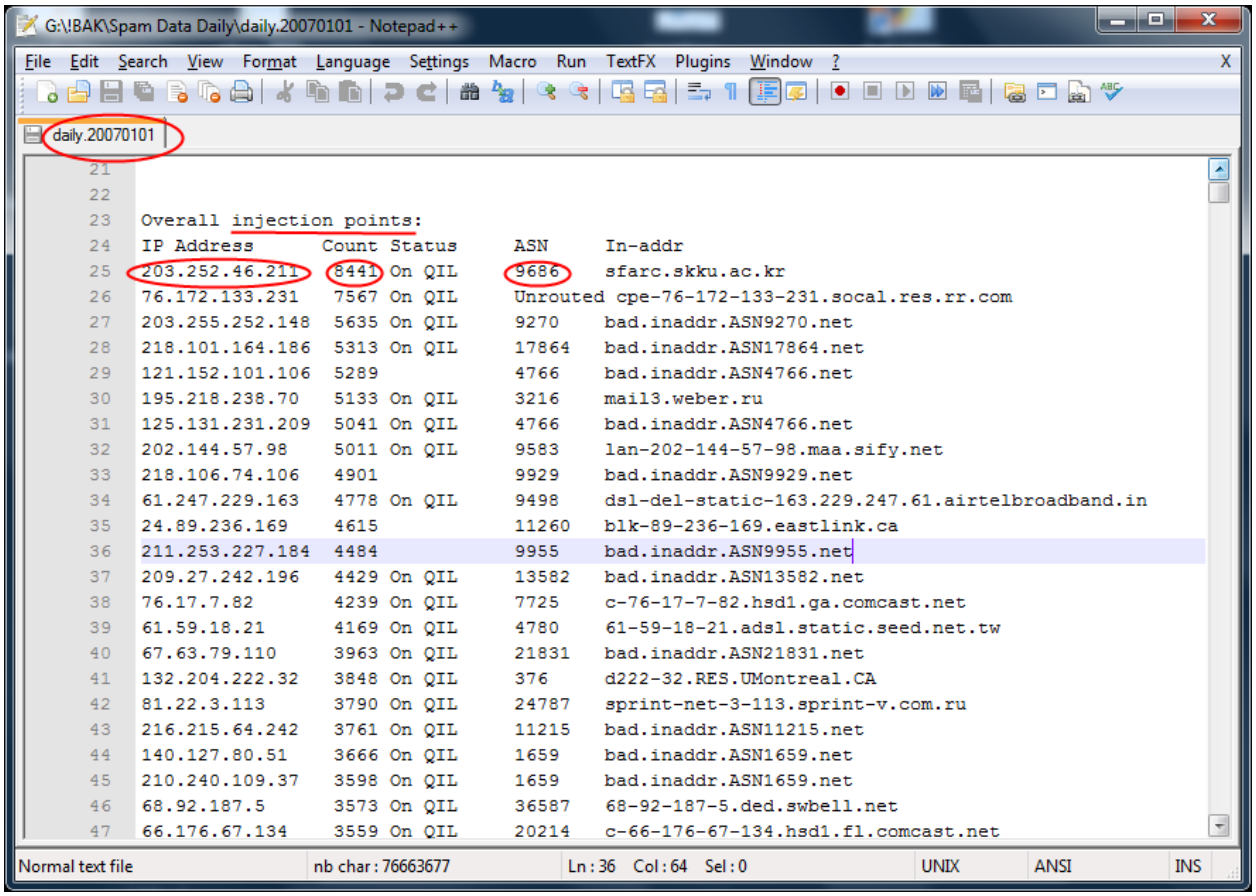

Figure 25 - Format of the raw spam logs

Now, this data is rather fine grained, as it has a record for each host emitting spam, where as we are interested in the aggregated ISP level activity. This is where the *Autonomous System Number* (ASN) comes into play. We will explain what the ASN is in more detail later in this chapter, but for now it's suffice to say that this number is an identifier for the ISP. Hence, using this field, and with the algorithm presented in Listing 1, we could extract ISP level metrics.

[1] Two important details have been left out for simplicity. The first is for each day there are two lists, called logged points and injection points, which for our purpose are the same. Second, the count field has a certain multiplier, different for each day. A brief explanation is given in the comments of the scripts presented in chapter 4.

Listing 1 – General algorithm for generating per ISP metrics

```
1. Create a set consisting of ASNs (≈ISPs) in memory, and for each ASN hold the
   following structure:

      [
      list of IP sources emitting spam from that ASN;
      total number of spam messages emitted from that ASN;
      ]


2. Loop through each line of each spam-log file (for the time period we're
   interested in), and do the following:

      i.   Extract the data fields from that line,
      ii.  Retrieve the structure of the relevant ASN from the set
      ii.  Add the IP address to that ASN's list of spam sources
      iii. Update the total number of messages sent from that ASN


3. Finally, create an output Excel file, where each row includes the following:

      (ASN, number of unique spam sources, total number of messages)
```

This algorithm will produce the first two variables, which we shall call ***unq_srcs*** and ***spam_msgs*** in the rest of this thesis. As you can guess, executing the above algorithm to produce the aggregated outputs is not possible by hand, and requires computer programming. For this purpose, we have opted to use the *Python* scripting language. We will look into the reasons why Python is the optimal choice for this task, and some explanations on the actual scripts in the next chapter.

A typical run of the scripts on the faculty's *'High Performance Cluster'* takes about 8 hours. The result is a dataset of several hundred thousand records from the nearly one billion log records. This dataset is again reduced when combining it with the independent variables to several hundred records.

Calculating relative performance

We mentioned the *'ceteris paribus'* clause in the examples we gave for our metrics. The reality is that when we compare any two ISPs, all else will not be equal. The most obvious difference is the size of the ISPs. As a result, ranking based on the metrics *unq_srcs* and *spam_msgs* will not necessarily identify the more secure ISP. To illustrate this point, consider the following two ISPs:

- ISP A, with 1,000 subscribers, 50 of which are infected;
- ISP B, with 1,000,000 users, and 2,000 infections.

In absolute numbers, ISP-A is better, where as it is rather obvious that ISP-B is actually doing much better (the infection ratio is under 1 percent versus 5 percent). Of course, many other factors can also influence the interpretation of these metrics (for instance, what if ISP-A provides ultra-high speed broadband, where as ISP-B provides dialup Internet?) [1] However, size is without doubt the most important factor, as ISP sizes range from a few thousand subscribers (small ISPs) to several million subscribers (mega ISPs). For this reason, we propose two other dependent variables, ***unq_srcs_sub*** and ***spam_msgs_sub***. These metrics are the absolute numbers divided by the number of subscribers of the ISPs.

Persistence

The creation of the third variable, the average duration of infections, involves more work. We call this metric ***persistence***. Calculating it involves tracking the number of days each IP is active in the algorithm presented above, and calculating the average in step 3. It also involves answering two questions:

- Will we look at the persistence of all machines, or just the worst offenders? (Who are worst offenders?)
- If a machine is seen spam on a Monday and then on the Thursday of the same week, will we consider the machine to have been infected on Tuesday and Wednesday as well? (Or was the machine disinfected and reinfected?)

'Persistence' promises to reveal interesting dimensions of an ISP's security performance. However, after some initial tests with the metric, it was concluded that refining it would involve considerable work, and hence it is not further developed in this thesis. Developing this metric can be a topic for further research.

[1] The reality is that all the other factors in the conceptual model must be controlled for to make truly accurate comparisons.

3.2.2 FACTORS THAT DISTORT THE RELIABILITY OF THE MEASUREMENTS

Several technical choices of ISPs can severely influence the dependent variables, which we shall review in this section.

DYNAMIC IP ADDRESS ASSIGNMENT WITH SHORT LEASE TIMES

Every computer connected to the Internet requires having an IP address, in order to communicate with other network hosts. In the assignment of IP addresses to subscribers, ISPs adopt one of the following policies:

- Assign a 'static' IP address to each subscriber. In this case, the IP address uniquely identifies that subscriber.
- Assign 'dynamic' IP addresses with long 'lease times': The subscriber receives an IP address that in theory can change, but, the ISP guarantees that the address will not be changed for a certain period of time (e.g., 30 days). Effectively, this policy is similar to the previous.
- Assign dynamic IP addresses with short lease times. In this scenario, each time the subscribers turns on her computer and connects to the ISP (or say every 24 hours), her IP address changes.[1]

Effects on the measurement tool

Technically these policies do not make much difference for ordinary users. For our measurement however, the third policy makes a big difference, and significantly over reports *unq_srcs.* To understand why, assume that an ISP has only one infected subscriber. In the case of the first policy, our logs will show *unq_srcs = 1* for a one month period. In the case of the third policy, the IP address of the infected subscriber changes, and we might see something as high as *unq_srcs = 30*!

Dynamic IP assignment doesn't affect *spam_msgs.* Regarding the distortion of *unq_srcs,* no clear remedy exists.[2] We simply need to accept the distortions caused by dynamic IP assignments as a limitation for our measurement tool. Among the security community, and in similar types of research, this limitation is considered acceptable, and IP addresses are generally used as proxies for users. (This can be confirmed by looking at how security sites such as SANS or the Conficker-Working-Group report everything in IP addresses; so do many academic papers focused on these subjects.) One reason is that because often no better proxy exists. [3]

NETWORK ADDRESS TRANSLATION

Network Address Translation, or NAT, is a technique used to share IP addresses between computers. Figure 26 shows how the concept works: the two computers on the left have what is known as a *private* IP address[4]. Private IP addresses are not 'routable' on the public Internet – that is, packets to and from these sources are discarded on the Internet backbones. The NAT gateway (which in many cases is a home router), 'translates' outbound packets by putting its own IP address as the source of the packet. When it receives the reply packets from the remote

[1] There are several reasons why an ISP would adopt such a policy: IP address sharing; selling static IPs at an extra price; regulations; etc.

[2] It seems that ISPs in the same geographic location adopt similar IP assignment policies, and as broadband penetration increases, so does the tendency to give static IPs to customers. (This certainly seems to be the case in the Netherlands).

[3] A mathematical solution comes to mind for detecting this anomaly: comparing the average number of *unq_srcs* per day, with the number of *unq_srcs* for the period divided by the number of days in the period.

[4] Private IP addresses are several address ranges that have been put aside by the Internet Assigned Numbers Authority specifically for the purpose of being used on internal networks.

server, it does the reverse.[1] Using NAT is common practice for ISP subscribers that have a small home (or office) network.

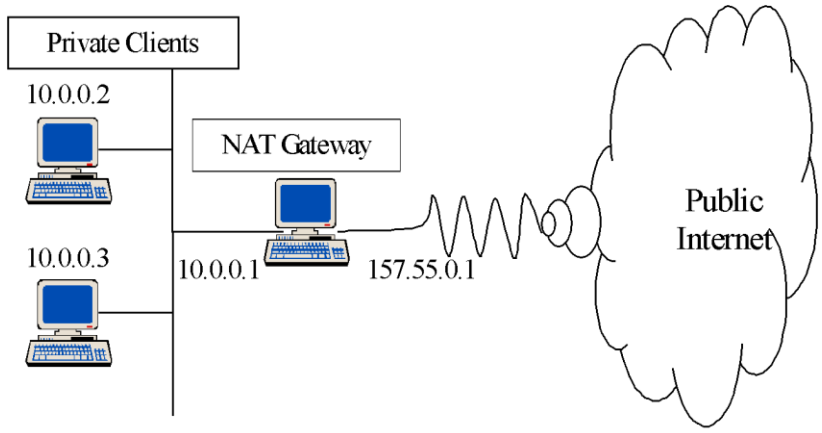

Figure 26 - Network address translation (image source unknown)

Effects on the measurement tool

But how does NAT influence our dependent variables? In the cases were multiple computers behind a NAT are infected, they will show up in our logs as one unique source. How big of a problem this would be is debatable. The argument goes that in the end, all the infected computers behind a NAT gateway belong to the same subscriber, usually have similar security conditions, and share the Internet bandwidth. Consequently, they can be accounted for as a single infection. The key thing to remember is that *unq_srcs* is more representative of the number of infected subscribers then infected PCs; and again, this is a generally accepted limitation in similar research.[2]

A much more serious problem happens if *an ISP does NATting*. This is, the ISP assigns private IP ranges to all subscribers, and NATs everyone behind a very big NAT gateway situated inside the ISP. Now, all the infected subscribers will appear as one source in our data. The ISP will appear to be ultra secure (e.g., *unq_srcs = 1*) were as in reality hundreds of thousands of subscribers might be infected! Luckily, NATting all customers in such a way is not recommended and in networks providing broadband access rarely occurs. [***] (It occurs mostly in developing countries where ISPs have a small IP address blocks)

Network address translation does not affect *spam_msgs,* as the spam-bots will happily continue sending the same amount of spam irrespective of this issue. [3]

OUTBOUND PORT 25 BLOCKING

One of the most recommended anti-spam measures in anti-spam best practice guides is *blocking TCP port 25*. In this technique, which is an attempt to stop bots sending out spam (and to close loopholes in the original Simple Mail Transfer Protocol), all outgoing traffic from this port is blocked for residential end-users. [4] This port number is reserved for mail servers, and since residential users do not normally run one, it's a safe bet any that outbound

[1] For this purpose, the NAT gateway keeps a table in memory of the translations it is currently doing

[2] As already mentioned, the claim can be verified by looking at security sites or scientific papers in Internet security.

[3] This can actually be used to detect ISP level NATs – having extremely unusual ratio of *spam_msgs* to *unq_srcs*

[4] Briefly put, the short comings in SMTP are *lack of authentication*, and *overloading of different responsibilities* in the protocol. Authentication is not implemented for *Mail Transfer Agents* (mail servers) – that is, a spammer can pretend to be sending legitimate email from any domain - without needing to prove ownership of the domain to the recipient MTAs. Overloading of responsibility means that stopping clients to act as MTAs, in order to solve this issue, distorts other aspects s of the protocol.

traffic using this port is a spam bot. As such, blocking the port would disarm the spam-bots. All valid emails are expected to be sent through the ISP mail relays, as shown in Figure 27.

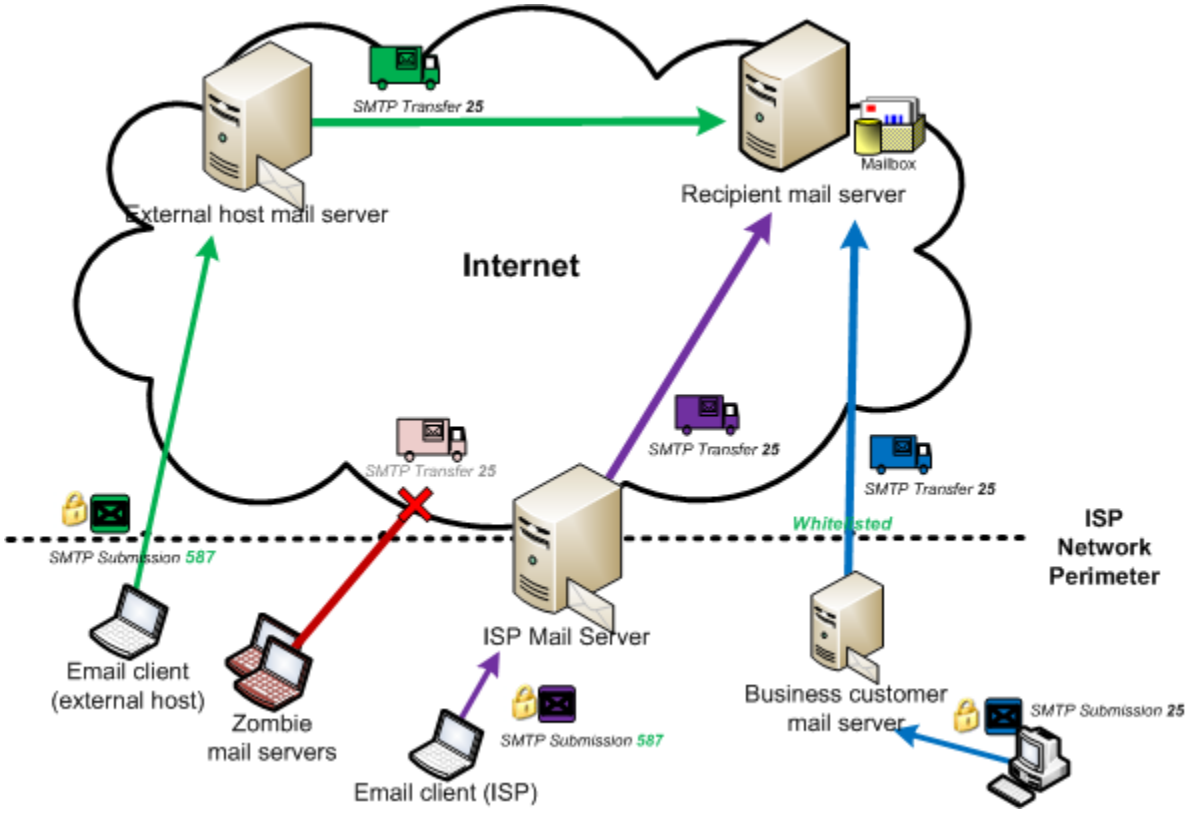

Figure 27 - Status of SMTP traffic after port 25 blocking has been put in place and legitimate customers have been either white-listed or advised to use the 'submission' protocol.

There is great controversy surrounding port 25 blocking (e.g., see comments on Kerbs 2009). Some of the arguments given against this measure as are follows:

1. The business side of ISPs are many times against implementing port 25 blocking, since it would stop smaller business customers from running their own mail-servers, and additionally, during the transition period that the measure is put in place, create disruptions for many customers.

2. The actual security benefits of the measure are also in doubt, as the determined criminals have already adapted their spam-bots to send spam via the ISPs relays. (In other words, for some ISPs that implemented this feature, only temporary drops in spam volumes occurred). Additional measures, such as outbound content filtering or traffic throttling on the mail relays need to be coupled with this measure to make it effective, which ISPs are reluctant to do.

3. Port 25 blocking disrupts monitoring systems that use spam as a proxy for botnet activity – including our dependent variable *unq_srcs,* and to a certain extent, *spam_msgs.* Some experts argue that this is like sweeping the dust under the rug - and destroying visibility of the bots, without actually remediating them.

Nonetheless, some other experts believe that port 25 blocking cuts into heart of one of the botnet business (spamming), if implemented together with the complimentary measures, and some ISPs have reported great success in keeping both the Internet clean and lowering their abuse levels (source: Dave Rand).

Effects on the measurement tool

How can we mend the problems created by port 25 blocking for our measurement tool? Foremost, we can detect ISPs that have implemented the measure incompletely. For these ISPs, *unq_srcs* drops considerably, but *spam_msgs* returns to approximately the same levels after an initial drop. One solution could then be to remove these networks from our dataset. For the group of ISPs that have fully implemented the measure (and thus have both metrics low, and a higher security ranking), one could argue that they are indeed being good citizens and deserve the higher rankings. In either case, this is another limitation we need to accept, and take precautions when interpreting our final results.

INSTANCES WHERE SPAM_MSGS GETS DISTORTED

In all three situations explored so far, the unq_srcs metric was distorted but the effect on spam_msgs was minimal. This might lead us to believe that spam_msgs is a more robust metric, but this is actually not true. The amount of spam emitted is influenced by, among others, the access speed of the subscriber.

Consider two ISPs with the same number of active spam bots. The one that has a higher average bandwidth per subscriber, or whose users spend more time online, would have a higher spam volume (*spam_msgs*). This is simply because each spam-bot can pump out more spam. Remember that our goal is to assess the differences between ISPs in terms of security effectiveness. The comparison of *spam_msgs_sub* becomes misleading in such cases.

Another example could be ISPs that offer shared webhosting services. If the web-hosts are hit by malware, they can emit enormous amounts of malicious traffic – after all, they are online 24 hours a day, and have very fast Internet connections. In such incidents, *spam_msgs* would greatly increase, but *unq_srcs* not, as the number of web-hosting machines is rather low compared to the total number of subscriber machines.[1] Since the goal is to assess botnet activity levels among subscribers, *unq_srcs* would be a truer reflection.

A related argument can be made regarding the interpretation of the actual value of the metrics. The *unq_srcs_sub* metric can be somewhat crudely interpreted as the percentage of computers in a network infected during a specific period. Interpreting the value of *spam_msgs_sub* is not as intuitive. This argument can also be seen as a plus for the *unq_srcs_sub* metric.

CONCLUSION

Throughout this section we investigated a battery of dependent variables extractable from our dataset of spam emails. We introduced three variables: *unq_srcs*, *spam_msgs*, and *persistence*, and also the normalized versions of the first two (*unq_srcs_sub* and *spam_msgs_sub*). These different variables measure different aspects of security effectiveness of an ISP, and as such are effected differently by the measures and policies of an ISP. We summarize all of the discussions of this section in the following table.

[1] For instance, 200 servers compared to 500,000 broadband subscribers

Table 13 - Comparison of the dependent variables

Dependent variable	Description	What security measures effect it[1]	What factors distort it
unq_srcs_sub	Number of unique IP sources emitting spam from a network during a specific time period. Corrected for size.	This variable indicates the effectiveness of measures taken before a computer is compromised. • Filtering (not outbound) • User education(pre inf., not post) • Client security (not quarantining) • Proactive detection (of new threats, not existing bots) • Participation in security community (not feedback loop)	• NAT (ISP level) • Port 25 blocking • Dynamic IPs
spam_msgs_sub	Total number of spam messages emitted from a network during a specific time period. Corrected for size.	This variable indicates effectiveness of both stopping machines being compromised in the first place, and speed of remediation afterwards. • Active abuse desk • Filtering • User education (pre inf. & post) • Client security (including quar.) • Proactive detection • Participation in community (including feedback loops)	• Port 25 blocking • Webhosting and similar services • Access speeds very different from the norm
persistence ***(not used)***	Average duration hosts remain infected in a network. This can be calculated for all hosts or a subset of them.	Similar to *spam_msgs_sub* – both measures that decrease infection rates, and increase the remediation process, affect this variable.	-

In the end, the key question remains: which dependent variable should we use? The two main candidates (*unq_srcs_sub* and *spam_msgs_sub*) both have their separate shortcomings. [2] It seems that it would be best to run the statistical analysis for both metrics. For a hypothesis to be accepted or rejected, we would require similar accept / reject results for statistical tests, when either of the metrics is used as the dependent variable.

Additionally, as we have said several times in this chapter, some of these shortcomings must just be accepted as limitations to our work. They could be mitigated in future research by adding other sources of security metrics. Luckily, these limitations are considered normal for the field, and as previously mentioned, in a workshop held in the TU Delft with security experts from the Dutch ISP XS4ALL and the anti-spam firm Cloudmark in September 2009, the method of using outbound spam (sources and volume) as a proxy for botnet activity was deemed reliable and valid.

[1] Management & administrative processes affect all three variables. The effect of legal measures and using updated servers and protocols is limited on bot infections.

[2] Some mitigation techniques for suggested in this section for resolving the distortions. However they were not pursued as they seemed both tricky and unreliable.

3.2.3 INDEPENDENT VARIABLES

Going back to our conceptual framework (presented in the previous chapter), we would need to add a variety of independent variables to be able to answer our research questions. Obtaining all these variables is not straight forward, and they can basically come from two sources: performing a survey on ISPs, or combining available secondary data sources. Taking into account the scope of this master thesis and the amount of time required to design an acceptable questionnaire and sending it out to several hundred international firms, we decided not to take this route, and instead opted for using secondary data from databases that we could gain access to.

TELEGEOGRAPHY

TeleGeography is a company that gathers and compiles statistics on the global telecom market. One of its commercial databases, *'GlobalComms'*, will be used as one of our main sources of data on Internet service providers. GlobalComms contains market data on wireline, wireless and broadband competition (TeleGeography 2009). The relevant part for us is the data on the 370 broadband service providers that they track. The data consists of a quantitative part which includes the following:

- Number of subscribers of the broadband providers, together with the type of service (DSL or cable)
- Financial information (such as revenue, CAPEX, and EBITDA margin) for a subset of these companies

The database also includes qualitative background information on each company and the market it operates in. This qualitative information can be useful for looking at certain companies in detail after the quantitative analysis, should the need for further clarification arise.

COUNTRY LEVEL DATA SOURCES

Several of the variables in the conceptual framework are influenced by country level variables – whether in a direct manner (such as laws and regulations), or in an indirect manner (the effects of country demographics on end user behaviour). Consequently, adding country level variables to our dataset can be fruitful. Obviously, since our unit of analysis is at the ISP level, only variables that influence the ISPs will be helpful.

World Development Indicators

The *World Bank Institute* compiles a widely endorsed database of over 700 country level indicators for the globe, called the *"World Development Indicators"*.[1] From them the following could be useful for our purpose:

- *International Internet bandwidth (bits per person):* possible link to the average ISP bandwidth
- *Broadband subscribers:* can be used to calculate the market share of each ISP, which could possibly be linked to certain incentives

Some other variables were also originally chosen but later dropped.[2] [3]

[1] The complete list of these indicators is available online at: *http://publications.worldbank.org/WDI/indicators*.

[2] The most important of these are: i) *GDP per capita* – a very interesting variable at the country level, but at the ISP level it is not useful as it is correlated with too many factors (e.g., it is an indirect indicator of end-user education, technical infrastructure, higher salaries for customer support, etc.), causing problems such as multicollinearity. ii) *Price basket for Internet* – could be linked to ISP incentives, but it is unfortunately not available for all years.

[3] Also please note that two other datasets were scanned for useful variables. One is provided by the ITU (key statistics on global ICT), and the other by OECD. No particularly exciting variable was found for our purpose in these datasets.

Cyber law regulatory index

A 'cyber law regulatory index' dataset has been compiled by our colleague Shirin Tabatabaie of TPM faculty. The dataset contains a variable indicating whether a country has signed and / or ratified a number of important cyber-crime laws, including the Convention on Cybercrime, and the London Action Plan. We will use this index to investigate the effect of country laws on ISP security effectiveness.

The purpose of the *London Action Plan* is to promote international spam enforcement cooperation and address spam related problems, such as online fraud and deception, phishing, and dissemination of viruses. The *Convention on Cybercrime* is a more general treaty that seeks to address computer crime and Internet crimes by harmonizing national laws, improving investigative techniques and increasing cooperation among nations. It covers much more than spam (e.g., illegal access, illegal interference, child pornography, etc), and falls under criminal law. (Please see the appendices for more information).

Global piracy rates

There is a widely held belief that software piracy causes cyber-insecurity. The reasoning is that many pirated software packages cannot be updated with security patches, and even more, some pirated software packages are infected with malware out of the box. (In our conceptual framework, software piracy rates would act as a technical enabler). The Business Software Alliance (BSA) published annual statistics on global piracy rates. The variable would be an interesting addition to our dataset.

Human Development Indices

The UNDP Human Development Reports provide a wide range of indices measuring 'human development' in different countries. Among these indices, *education index* can be a useful variable, as it measures the education level of an ISP's end-users (based on the country average of course), and this indirectly and crudely can be linked to the online awareness of those users.

OTHER POSSIBLE SOURCES OF INFORMATION

It might be possible to extract certain variables manually from the following sources:

- Company annual reports (e.g., R&D expenditure as a measure of company culture);
- Company web sites (e.g., regarding services and policies);
- ISP industry associations (regarding member firms);

Such additions are however not planned at this stage, as they will be to labour intensive.

3.2.4 COMBINING THE VARIABLES INTO ONE DATASET

At this point we have identified the dependent and independent variables that we wish to use for our statistical analysis. Combining these variables into one dataset has some challenges of its own. Combining the country level independent variables and matching them to the broadband providers listed in TeleGeography is somewhat time consuming, but not particularly tricky. However, coupling them with the dependent variables is far from trivial. To understand why, we need to recall the format of the spam logs (see Figure 28). As we explained in section 3.2.1, the logs are at the level of the individual spam sending machines, which would have to be aggregated to the ASN level. We briefly mentioned that ASNs can be mapped to ISP names. This task is by no means straightforward, as will be made clear after we explain what an ASN exactly is.

```
23  Overall injection points:
24  IP Address       Count Status    ASN     In-addr
25  203.252.46.211   8441 On QIL     9686    sfarc.skku.ac.kr
26  76.172.133.231   7567 On QIL     Unrouted cpe-76-172-133-231.socal.res.rr.com
27  203.255.252.148  5635 On QIL     9270    bad.inaddr.ASN9270.net
```

Figure 28 - Excerpt of the raw spam logs

AUTONOMOUS SYSTEM NUMBERS

A premier on Internet routing

Internet routing is architectured around the concept of *'Autonomous Systems'*[1]. The Internet is split into a large number of autonomous systems, such as ISPs, large corporations, and universities. Each autonomous system controls a range of IP addresses, and is responsible for its internal routing policies. Each autonomous is assigned a unique number, called the autonomous system number. Figure 29 shows some example ASNs.

Figure 29 - Example ASNs

On the boundaries of each autonomous system, border 'routers' exchange reachability information about the address blocks they control with other autonomous systems. (The Border Gateway Protocol version 4 is used for this purpose). Figure 30 shows how this works – here, AS 1 is advertising an IP address block starting with 11.11.*.*. This information gets propagated throughout the Internet. A computer residing in AS 19222 will have two possible paths to choose among, should it wish to send a packet to the specified address block. (Most likely the shortest path will be chosen).

Autonomous systems taking part in global Internet routing receive their ASN from one of the five Regional Internet Registrars (RIPE in Europe, ARIN in the US, and APNIC, LACNIC, and AfriNIC in other parts of the world.)[2] The range of ASNs is between 1 and approximately 65,000; today more than half of these numbers are in active use.

The first step in mapping an IP address sending spam to its ISP would be to look at the Internet routing tables at any point in time and finding the AS advertising the block that IP belongs to. Luckily, this has already been done in our data, and the result has been saved in the ASN field. The next step is to find which ISP owns this ASN.

[1] Routing is the process of selecting paths in a network, along which to send traffic (i.e., when your computer connects to a website, what networks do the data packets traverse in order to reach the remove server?).

[2] The Internet Assigned Numbers Authority coordinates the ASNs and IP address blocks between the RIRs.

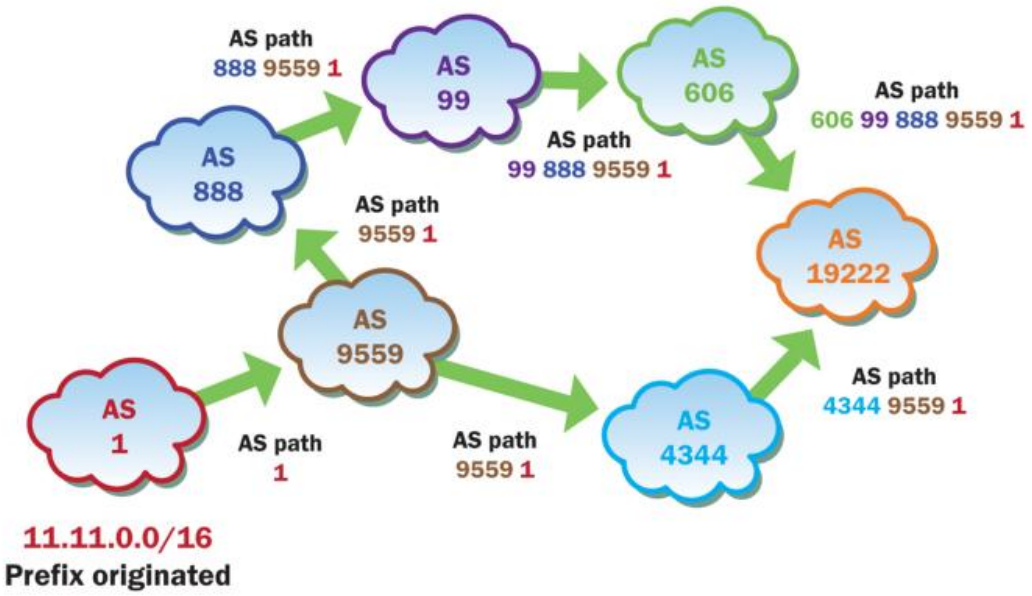

Figure 30 - Path propagation in BGP (source: renesys.com)

Mapping ASNs to ISPs

Using data from the sources such as Route Views[1], we can look up the name assigned to the autonomous system by its owner. In some cases, we can derive the name of the ISP from the AS name. However, several complications exist. One is that due to mergers and acquisitions, ASNs change hands, yet in the process, retain their original name. (The parent company might decide to merge the new AS into its existing network sometime in the future). As an example, Table 14 lists the major ASNs owned by the Dutch incumbent telecom, KPN, and demonstrates this point.

Table 14 - List of majors ASNs owned by KPN in 2007

ASN	Name	Usage / History
AS286	KPN KPN Internet Backbone AS	KPN's core network, European backbone, and some business customers are on this AS
AS3265	XS4ALL-NL XS4ALL	The daughter company XS4ALL retains a separate ASN
AS5417	DEMON-NL Demon Netherlands, TDINL BV	Demon was acquired by XS4ALL. This AS was later merged into AS3265.
AS5615	TISNL-BACKBONE Telfort B.V.	Tiscali was acquired by KPN, and later the brand was renamed to Telfort.
AS8737	PT KPN Internet Solutions	Residential KPN subscribers are on this separate AS (customers of brands such as Het Net and Planet)

Unfortunately, no automated method exists to identify the ASNs that belong to a particular ISP. The mapping has to be done manually, by Googling the AS names, and looking for clues as to which of the broadband operators named in the TeleGeography database it should be linked with. This process is rather labour intensive. More information on the execution of this procedure will be given in chapter 4.

[1] http://www.routeviews.org/

IP Location

Another complication arises from the fact that autonomous systems are not confined to country borders. For instance, the cable company UPC, owns AS6830. This AS runs across the Netherlands, Hungary, Austria, Czech Republic, Ireland, Slovakia, and Belgium! In this case, which of the country level variables should we match with this ASN in our dataset?

This is where *IP location* databases (also known as *GeoIP*) come into play. Our solution is to query each IP address in the spam logs against an IP location database, and actually aggregate the data at an ASN/Country level. As an example, the final UPC data for 2008 is presented in Table 15.

Table 15 - Breakup of AS6830 (UPC) data across country borders

ASN/CC	unq_srcs	spam_msgs
6830/NL	165,190	120,608,288
6830/HU	161,748	111,601,853
6830/AT	90,026	67,341,122
6830/CZ	41,646	35,445,649
6830/IE	13,183	19,406,191
6830/SK	13,481	11,134,725
6830/BE	24,907	6,024,766
6830/FR	964	831,539
6830/other	2,809	742,091

IP location is not an exact science, and will never be 100 percent accurate. This is mainly because of the way IP blocks are handed out. The Regional Internet Registrars (who are responsible for the handouts), keep '*WHOIS*' data on the IP blocks; however the WHOIS database only keeps the address of the organization who the block has been registered to, not the actual location that the block is being used. (For instance, address blocks belonging to AOL are registered in the US, but AOL might have actually assigned them to some computers in Germany).

The creators of IP location databases combine multiple sources of data (in addition to the WHOIS data) to determine the precise geographic location of a particular IP address. Many times these databases do not agree with each other. We have chosen to use the *MaxMind GeoIP* database, which holds a certain *de facto* status.[1] It should be noted that GeoIP inaccuracies are an inherent limitation of our research and similar projects.

[1] At one point we performed a literature search to see what IP location database has been used in similar research, and found that MaxMind is overwhelmingly the most popular. One reason is that MaxMind offers a free version of their database (called GeoLite), which although slightly less accurate than their commercial GeoIP database, is used in many open source applications. For optimal accuracy, we have used the commercial version of their product.

THE FINAL OUTCOME

Putting together all the variables discussed in the proceeding sections, we can present our final dataset as follows:

Table 16 - List of variables in our final dataset

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 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	
	spam_msgs_sub	Spam messages per subscriber. Similar to spam_msgs, but corrected for size	
Independent variables	*total_subs*	Total number of subscribers of an ISP (retail, business, DSL, cable, etc)	TeleGeography
	srv_type	The type of service / access provided by the ISP: DSL, cable, or both.	
	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).	WDI
	bb_subs	Number of broadband Internet subscribers in the country the ISP operates in. (note: we use this variable indirectly, in calculations)	
	lap_mem	Is the country of the ISP, a member of the London Action Plan?	Mother OECD project
	cyberconv_mem	Has the country of the ISP, signed the convention on cybercrime?	
	piracy_rate	Percentage of software that is pirated in the country the ISP operates in.	BSA
	educ_ix	Education index: an index indicating the overall education level of people in the country that the ISP operates in.	UN HDR
	market_share	Local market share of the ISP (total_sub divided by bb_subs)	-
Mappings	*ASN-to-AS-name*	Mappings of Autonomous System numbers to names	RouteViews
	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

3.3 Formulating the Empirical Hypotheses

3.1.1 Constructing the Measurement Model

In the final section of this chapter, we will formulate the hypotheses that need to be tested. Ideally, we would like to test all the relations present in our conceptual framework, and build a full regression model. However, despite all our efforts in gathering and combining secondary data sources, we do not have the necessary data to perform extensive testing. Our dependent variable is rather rich, but unfortunately, our independent variables are somewhat restricted. This means that for this thesis, only part of the relations in the conceptual framework can be tested. (The positive side is that this leaves plenty of room for further research). In Figure 31 we show again the conceptual framework, highlighting the components that we have data on, and greying out those that we don't.

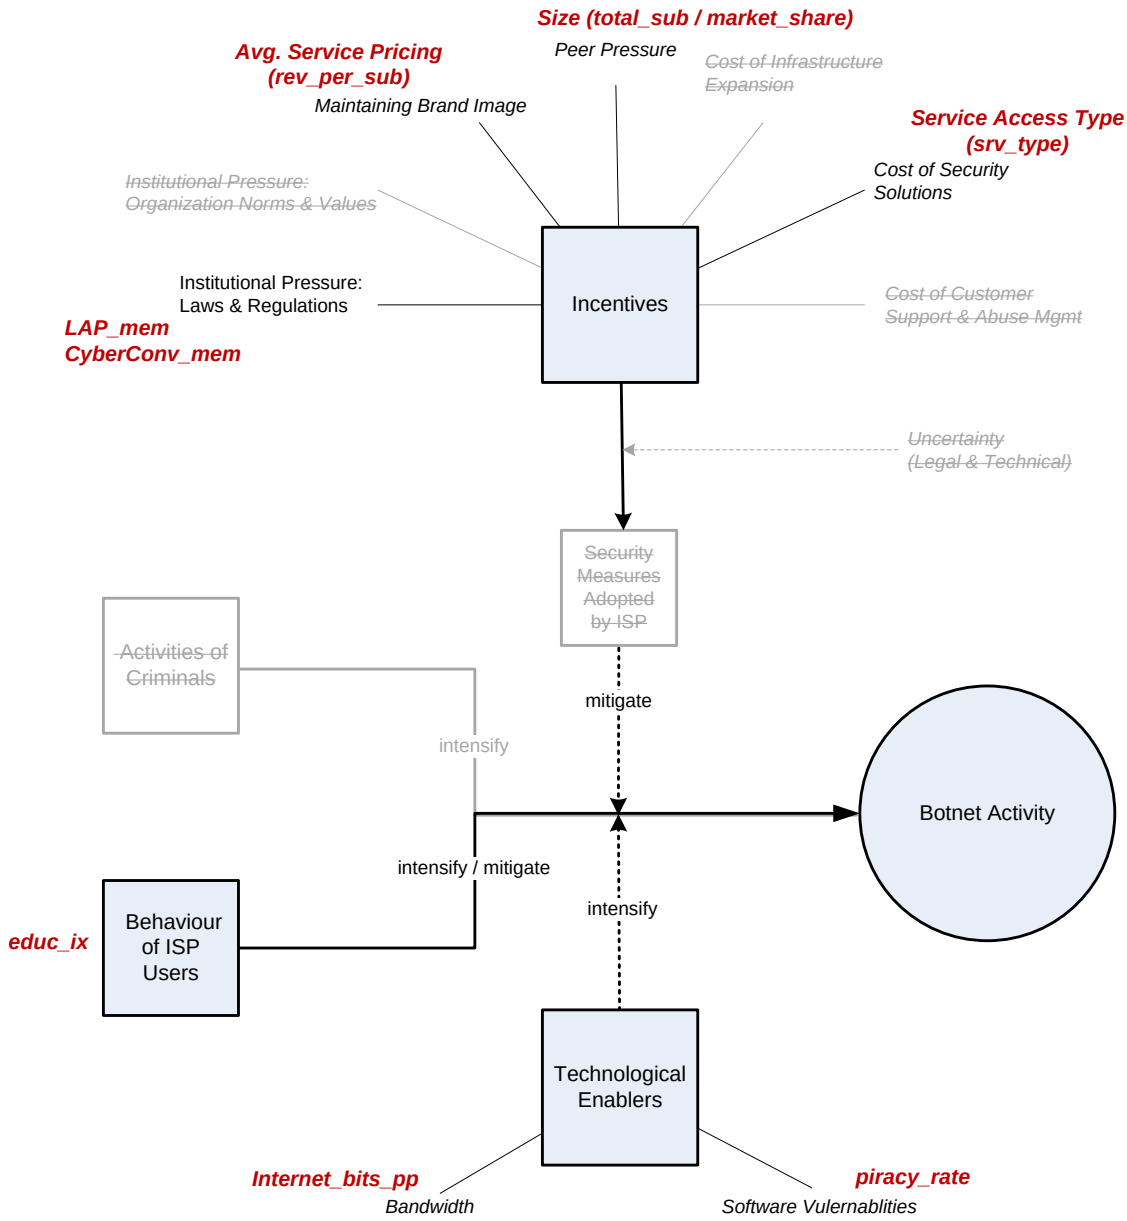

Figure 31 – Moving from the theoretical model towards a measurement model

The highlighting and greying has been done by taking the list of variables given in Table 16, and matching them with the terms in the framework. Please not that in many cases, these are rather crude proxies. For instance, whereas the number of subscribers of an ISP might be an acceptable proxy for its size, education index is a rather weak proxy for user awareness (of an ISP's users). Despite being aware of such problems, we opt to retain all the variables at hand in the model, as this broadens the number of hypothesis we can test. Consequently, care must be taken in interpreting and generalizing the results.

If we take out the greyed areas from the model, we will end up with our *'measurement model',* presented in Figure 32. In the figure, the elements are reordered but the grouping is kept similar to what we had in the conceptual framework (theoretical model). There is one major change however - the moderating relations have been changed to direct relations. The rationale behind this is that although our theoretical model suggests that botnet activity is caused by the end users and criminals, and moderated by ISP security measures, we practically have no measurements in our dataset for the causes, and hence, examining the effects of moderation on those links becomes impossible.[1] [2] For this reason, all the relations in the measurement model are simple, direct relations.

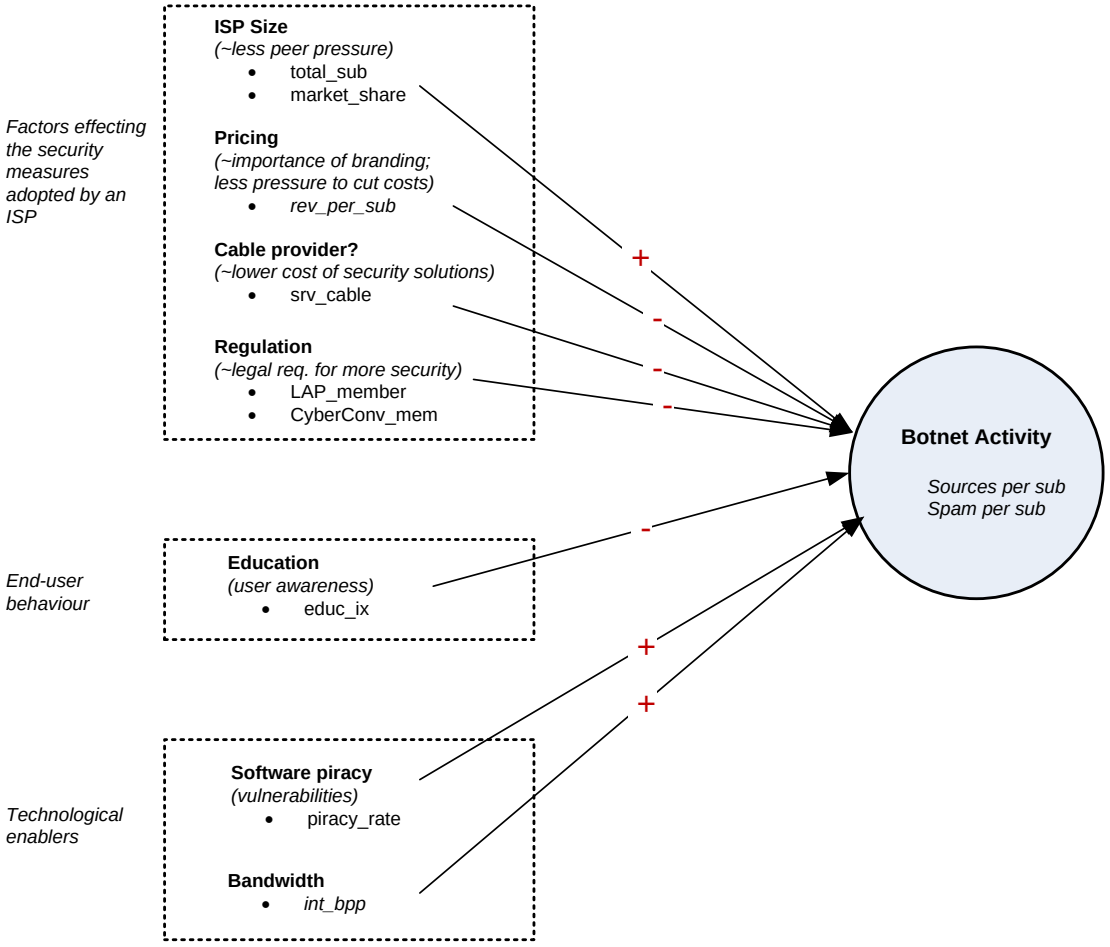

Figure 32 – The measurement model

[1] Please note that although 'criminal activity' was present in the conceptual framework, its absence in the measurement model is not problematic. This is because in the long run, we have no reason to believe that criminals would favour targeting a specific ISP (among ISPs with similar characteristics that is) more than its peers. In other words, in addition to being very hard to quantify, there is actually not a need to control for criminal behaviour, as it would already be captured in ISP characteristics.

[2] 'educ_ix' is too weak as the only indicator for end-user online behaviour, so it doesn't really count either.

3.1.2 LIST OF PROPOSED HYPOTHESES

Based on the measurement model, and our research questions, we can develop a set of empirical hypotheses. In this section, we present nine such hypotheses. The first two are related to sub-questions 2 & 3, and can be answered empirically with the help of the dependent variable alone. The next seven are based on the 7 relations in the measurement model, and jointly shed light on sub-question 4.

SubQ2: …. Are ISPs the key intermediary in botnet mitigation efforts?

Hypothesis 1: A majority of the world's malicious traffic originates from "autonomous systems" run and controlled by a limited number of ISPs, as compared to originating from autonomous systems controlled by other types of organizations (enterprises, universities, etc).

Explanation: *If this hypothesis holds, we can regard, ISPs as the "gatekeepers" of the botnet problem, i.e., they are crucial intermediaries in botnet mitigation, as a big chunk of the global botnet activity can be reduced by certain decisions of this group of actors.*

SubQ3: Do ISPs significantly differ in the degree in which they mitigate botnets?

Hypothesis 2: ISPs significantly differ in terms of their performance in mitigating botnets.

Rationale: *Due to the MIXED incentive structure and uncertainties, ISP's adopt very different security measures, which we expect to result in significantly different levels of botnet mitigation (as measured by the level of botnet activity.)*

SubQ4: To what extent can we explain the varying degree in which ISPs mitigate botnet activity? Can we identify internal or external factors that can explain this variance?

Hypothesis 3: Larger ISPs perform worse in terms of security (i.e., have lower security performance[1])

Rationale: *In the literature it is believed that having a larger size would lower peer pressure; and hence reduce incentives for security; (Also, bigger ISPs usually have over-capacity in bandwidth);*
Opposing argument: *reputation (branding) is probably more important for large companies (as they rely less on word of mouth), which is an incentive for security; also larger ISPs usually perform automation that will then lower the cost of security measures, so security increases*

Hypothesis 4: ISPs with higher average revenue per user have higher security performance

Rationale: *ISPs involved in price competition would worry less about brand image (reputation); they would also be forced to cut down on all costs, which include investing less in security measures;*
Opposing argument*: on the other hand, they might want to lower abuse management costs, and cost of infrastructure expansion, and actually taking a longer term perspective, increase their security measures.*

[1] Please note that in the text of the hypotheses, security performance is defined in terms of botnet activity. In other words, a lower security performance means higher levels of botnet activity; and similarly, higher security performance means lower levels of botnet activity – in relative numbers of course.

Hypothesis 5: Cable providers have higher security performance than DSL providers

Rationale: *Cable providers already have equipment in place to monitor network traffic (due to having a shared-network infrastructure). This same technology can be cheaply adapted for security purposes;*
Another argument is that since cable providers have mostly retail users, they can adopt simple, strict security policies (reducing the complexity, and hence costs of security measures).

Hypothesis 6: ISPs in countries that have endorsed international agreements against cyber-crime (e.g., the LAP or the cyber-crime convention) have higher security performance

Rationale: *The law requires them to be more stringent in terms of security*

Hypothesis 7: ISPs in countries with higher piracy rates have lower security performance

Rationale: *Using pirated software increases the number of exploitable software vulnerabilities, and is often mentioned as cause of decreased security.*

Hypothesis 8: ISPs in countries with higher average bandwidth rates have lower security performance

Rationale: *The average bandwidth of a country (i.e., broadband speeds) would be an approximate representative for the average bandwidth of subscribers of all ISPs in that country; since bandwidth is a technological enabler of botnet activity, this relation would hold*

Hypothesis 9: ISPs in countries with a higher educational index have higher security performance

Rationale: Education increases the awareness of an ISP's end-users in regards to Internet security risks (hopefully!), resulting in less risky online behavior of the users

(Please note that the term *security performance* is defined in terms of *botnet activity levels* – see footnote on previous page for more information).

Conclusion

In this chapter we defined our research methodology. We started by exploring various sources of data, and listing the dependent and independent variables available to us. Due to the limited number of variables in our dataset, testing the complete conceptual framework (theoretical model) is not possible, and hence a measurement model was constructed. With the help of this measurement model, a set of empirical hypotheses were developed, that when tested will answer our research questions.

CHAPTER 4- DATA PREPARATION

4.1 AGGREGATING RAW DATA USING PYTHON SCRIPTS

In the previous chapter we mentioned that our raw spam logs contain an enormous number of records (approximately one million IP addresses logged daily). While this level of fine grained data can be very valuable for certain tasks, for the statistical analysis that we seek, an aggregated dataset (at the level of ISPs, and on a yearly basis) is needed. The approach developed during this thesis (and the broader research it lies in) has been to write scripts specifically for the purpose of processing the raw log files, and compiling the eventual dataset, using the Python programming language.

4.1.1 BUT WHY PYTHON?

Python scripting is a very efficient tool for such scenarios. To understand why, let us consider some of the alternatives. The most straight forward approach that comes to mind is to load the gigantic raw dataset into a statistical program. This is simply is not doable – due to constraints of such programs (and even if it were, it wouldn't be of much practical use, since our units of analysis are much larger than individual IPs). A second approach that comes to mind is to load the data into a spreadsheet such as Microsoft Excel, and perform certain aggregation there. This is again not viable, as spreadsheets have extra bells and whistles, and typically try to load the whole data-file into memory at once. Both impractical, and an overkill for our purpose of executing the aggregation algorithm outlined in section 3.2.1.

A third option would be to import the logs into an RDBMS, such as MySQL. This option, although more practical than the previous two, is still not as efficient as we would like. For instance, databases typically create indices for each record, unneeded by us. They also add overhead to the data, increasing its size on the disk. And finally, database queries can take a very long time to execute on such number of records.

This is where Python scripts come into play. Python is fast in tasks such as processing text files, and executing an algorithm similar to the one we presented in section 3.2.1 (to recap, the algorithm was to go through all the log files, read the records, and update an aggregated version of the data held in memory. After the last file, the data structure in memory is saved to a CSV file). Compared to other programming languages, Python has unique advantages for the application we see:

- It has many useful data types built-in, such as dictionaries and lists, that speed up programming;
- A large collection of modules has been written for Python, many of which are open source, that add new functionality (e.g., for performing GeoIP lookups); these modules are in addition to Python's own excellent Standard library; together, they simplify many mundane programming tasks;
- Since Python is a scripting language, commands can be tested on the fly, memory handling is done automatically, and scripts can be easily 'stitched' together; all excellent features for rapid prototyping;
- And last, but definitely not least, Python is an easy to learn programming language, and enjoys having an elegant and clean syntax.

4.1.2 BUILDING THE PROCESSING INFRASTRUCTURE

It soon became evident that despite all the benefits of Python scripting, due to the large volume of data, and the number of different steps involved in the processing the logs, a computationally intensive task was at hand, requiring much horsepower and memory, and approximately 100GB of free disk space. These needs led us to use the ***High Performance Cluster***, a computing infrastructure setup for research involving intensive number crunching at TPM faculty.[1] Luckily, our requirements where easily met by this infrastructure.

A typical run of our longer scripts takes about 8 hours to execute for the full 2005-2008 range, and consumes around 6 GB of RAM[2]. This process is speeded up by running multiple instances of the script, in parallel and on different CPUs of the machine - each processing a different year. This lowers the execution time to around two hours. However, this is still not fast enough, due to the fact that for different formats of output, the scripts have to be modified and re-run. For instance, they would be run once to generate the list of the world's top spam sending countries; modified versions of the scripts would be rerun to graph the global spam trends; they would again be changed and rerun to produce the amount of spam sent by each ASN in the Netherlands, and so on. If you take into account bugs that can occur in the scripts (necessitating reruns), you can see that the solution soon loses its practicality. A speedier solution was needed.

Eventually, the hybrid solution presented in Figure 33 was developed. The raw logs are converted into an intermediate form (with one level of aggregation), and stored in a MySQL database[3]. To compile data into the necessary output formats, scripts and SQL queries are run against this database, with the runtime reduced to just a few minutes. Mapping ASNs to operator names, as well as joining the independent variables to the dataset is also done in this step.

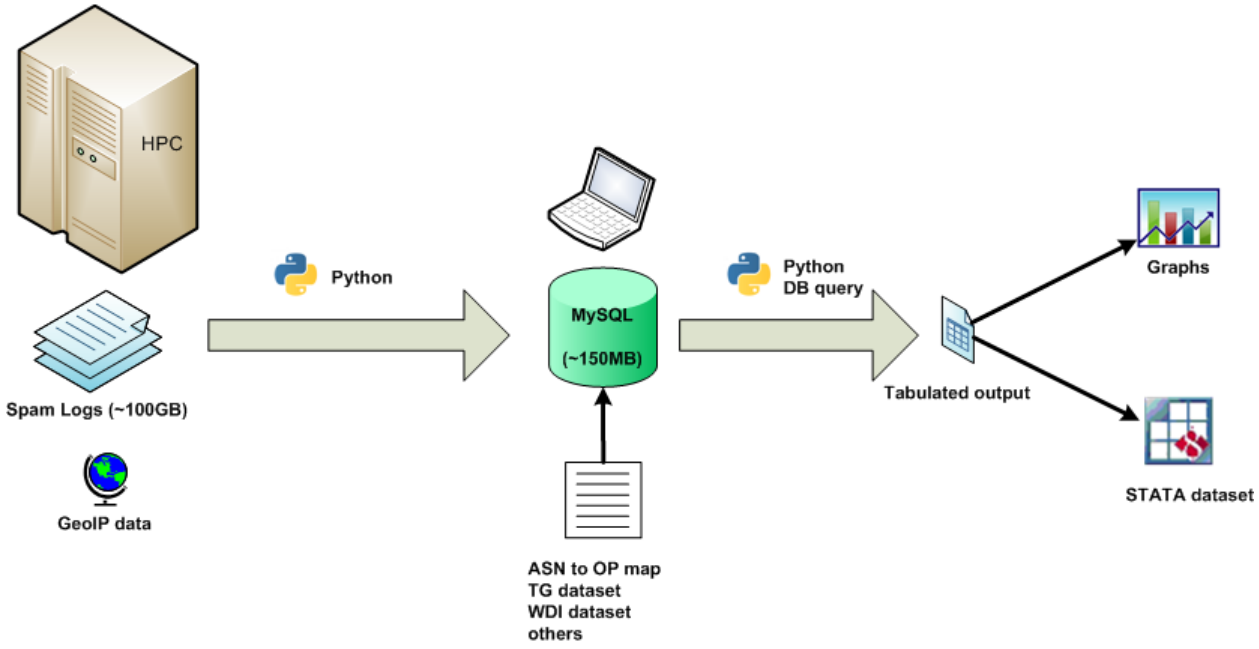

Figure 33 - Steps involved in compiling the final Stata dataset

[1] The HPC's specifications are as follows: 60 servers, each with 8 CPU cores, and 16 GBs of RAM, all running 64bit Linux. The cluster has access to a centralized 10 TB disk storage.

[2] Six Gigabytes seems to be a hard limit in the Python libraries, as after this an "out of memory" exception occurs.

[3] Time-wise, the data is aggregated at the weekly, monthly, quarterly, and yearly levels, and unit-wise, at the ASN/CC level.

4.1.3 SAMPLE SCRIPTS

Other research projects that involve processing large amounts of raw data into aggregated datasets can, similar to ours, benefit from Python scripting. For this reason, giving a more concrete example of the scripts can be academically valuable. In this section we will present two of the most important scripts used in the process. [1] The chosen scripts are from the set of scripts that process the raw log files. (Table 17 shows the various script categories.) The scripts are pretty self explanatory, and comments have been added to them where necessary.

Table 17 - Categories of the various Python scripts used in data prepration

Script functionality	Explanation
Check data consistency	Scripts that perform various consistency checks on the raw data
Process raw to spreadsheet (obsolete)	The initial set of processing scripts that directly transformed log files into the spreadsheet files and graphs. (Listing 2 is from this category)
Process raw to MySQL (newer)	The second set of processing scripts that transformed log files into the intermediate MySQL db. (Listing 3 is from this category)
Import other sources	Scripts that import the other data-sources (e.g., ASN mappings, TG, WDI, etc) into the database.
Compile Stata dataset	Scripts that produce the final Stata dataset from the database, performing various calculations along the way (e.g., calculating market share, growth, etc)
Concept tests	Scripts that perform tests on newer concepts (e.g., persistence, percentile ranking, etc).

The first script is rather simple, and involves aggregating the logs for one year, producing country level statistics, and saving the results to a file.

Listing 2 - Source code for generate_allcos_csv.py

```
# Script to generate spamcount/unqips for all countries for annual 200x
# Author Hadi. v1, 25 April 2009; v1.3 20 July 2009; simplified for print December 2009

import GeoIP
import os

yr = 2007                   # change this figure to the year that's logs are to be processed

country_unqips = dict()     # has items of form 'cc': set(ips) - holds IPs sending spam, per country
country_spamcount = dict() # has items of form 'cc': spam_count - spam volume, per country
daycount = 0                # number of log files processed
gi = None


# these outer loops execute once for each day of the year, (hence, each log file)
for mo in range (1,13):
    for dy in range (1,32):

        # load GeoIP database nearest to log-file date
        geoipdb = 'geoip/GeoIP-106 %4d%02d%02d.dat' % (yr, mo, dy)
        if os.path.isfile(geoipdb):
            gi = GeoIP.open(geoipdb, GeoIP.GEOIP_MEMORY_CACHE)

        # open the daily spam log file if exists for this day
        try:    fd = open("spam.daily/%4d/daily.%4d%02d%02d" % (yr, mo, dy) )
        except: continue

        # get 'avg' figure for this day. this is a multiplier for the actual spam volume
        while fd.next() != "Log file retry statistics\n":
            pass
        fd.next()
        s = fd.next().split(',')[2].strip()
        avg = float ( s[s.find(' '):] )
        if avg < 1 or avg > 15: avg = 4  # use 4 in case of unusual values
```

[1] These two scripts are chosen among more than 50 scripts written and used during this project.

```
        # now, go back to start of file to start processing records
        fd.seek(0)

        # log file consists of two parts, called 'injection points' and 'logged points'
        # the diff. being whether the spam source was blacklisted at moment of reception
        # for us they make no diff. and parsing is almost identical
        for section in ['injection', 'logged']:
            is_injection = (section == 'injection')

            # skip lines until section start
            section_start = "Overall %s points:\n" % section
            while fd.next() != section_start:
                pass
            fd.next()

            # this loop executes for each line until section finishes
            for s in fd:
                if s == '\n': break

                # parse line - the format is [ip count asn] - asn not used in this script
                ip = s[0:15].rstrip()
                count = int(s[16:21]) if is_injection else int(s[16:25])
                #asn = s[33:40] if is_injection else s[42:49]
                #asn = int(asn) if asn[0].isdigit() else -1

                cc = gi.country_code_by_addr(ip) or '--'         # perform GeoIP lookup!
                # if country not seen before, setup data structures
                if cc not in country_unqips:
                    country_unqips[cc] = set()
                    country_spamcount[cc] = 0

                country_unqips[cc].add(ip)          # add to list of ips emitting spam for country
                country_spamcount[cc] += int(round(count * avg)) # increase spam count for country
            #
        #
        fd.close()
    #
#

# section to save proccessed data in memory to a CSV file
fw = open('xspam' + year_s + 'global.csv', 'w')

header = 'CC, Spam Messages, Unique Sources\n'
fw.write(header)

for cc in sorted(country_unqips.keys()):
    row = "%s, %d, %d\n" % (cc, country_spamcount[cc], len(country_unqips[cc]))
    fw.write(row)

fw.write('Total days in period: %d\n' % daycount)
fw.close()

print 'Total days processed: ' + str(daycount)
```

The second script is much more complex. It aggregates the data to the ASN/CC level, than saves the results to MySQL, and prints its progress during run. The main loop is similar to the previous one.

Listing 3 – Source code for spam_by_asn_yr_mysql.py

```
# Script to generate annual spammsg_count and unqips_count, for all ASN/CCs.
# The result is inserted into MYSQL db.(Script uses daily log files)
# Author Hadi. v1, 21 July 2009, simplified for print December 2009

import GeoIP
import MySQLdb
import sys
import time
import os
import struct
import socket
from datetime import date


# this script uses a custom data structure to hold information for each ASN.
# that is, one instance of this class exists in memory for each ASN
# inside the class, spam_msg and unq_ip_src for that single ASN are held, divided by country
class AsnData:
    def __init__(self, n):
        self.myasn = n
        self.unqip bycc = dict()  # dictionary holds items of this form: 'cc' : set(ips)
        self.spam_bycc = dict()   # dictionary holds items of this form: 'cc': spam_count
        self.spam_total = 0       # grand total spam_count

    # this class-method is called to add one IP record
    def add_item(self, ip, cc, spam_msgs):
         # store IP addresses as integers - slow but conserves memory
        ip = struct.unpack('I',socket.inet_aton(ip))[0]
        self.spam_bycc[cc] = self.spam_bycc.get(cc,0) + spam_msgs       # update spam-count
        self.unqip bycc.setdefault(cc, set()).add(ip)                   # store IP
        self.spam total += spam msgs

    # this class-method is called write before outputing results to DB
    # to avoid db explosion, group all small ccs as other ('**')
    def group_smallcc(self):
        for cc in self.spam_bycc.keys():
            perc = 1.0 * self.spam_bycc[cc] / self.spam_total
            if perc < 0.001:
                self.spam bycc['**'] = self.spam bycc.get('**', 0) + self.spam bycc[cc]
                self.unqip bycc['**'] = self.unqip bycc.get('**', set()).union( self.unqip bycc[cc])
                del self.spam_bycc[cc]
                del self.unqip_bycc[cc]
        #
#

# global method that outputs results to MySQL
def write todb(a asnlist):
    print_flush( '   >> committing to db--%d asns @t:%.1f...' % (len(a_asnlist),time.time()-stt) )
    if len(a_asnlist) == 0:
        print '   >> 0 rows committed'
        return

    # loop until db becomes available, and open a connection to it
    print '   >> connecting to db...',
    while True:
        try:
            mydb = MySQLdb.connect(host='127.0.0.1',port=3366,user='XX',passwd='XX',db='spamdata')
            if mydb != None:
                print '   >> ... connected!'
                break
        except:
            print_flush ('error connecting to db, sleep 10 min....')
            time.sleep(600)
    #
    commit count = 0
    c = mydb.cursor()

    # loop over all the asns
    for asn, data in a_asnlist.iteritems():
        data.group_smallcc()
```

```
        # make sure that ASN-0 is really unused, because we are using 0 for unrouted/invalid asns
        if asn == 0: raise "we shouldn't have a zero asn!!!"
        if asn == K UNROUTED : asn = 0

        # loop over all ASN/CCs in this asn, committing one by one
        for cc, spam_count in data.spam_bycc.iteritems():
            ip_count = len(data.unqip_bycc[cc])
            sql = 'insert into spam_by_asn_yr (yr, asn, cc, spam_msgs, unq_ips) values ' \
                  '(%d, %d, "%s", %d, %d)' % (glbl_yr,asn,cc,spam_count, ip_days)
            try:
                c.execute(sql)
                commit_count += 1
            except MySQLdb.IntegrityError, e:
                print '*** IntegrityError: %s' % e # usually due to duplicate. see message details
            except Exception, e:
                print "*** Exception type %s: %s   for query '%s'" % (type(e),e , sql)
        #
    #
    c.close()
    mydb.commit()
    mydb.close()
    print flush('   >> %d rows committed  @t:%.1f' % (commit count , time.time()-stt) )
#

# this method flushes output after printing a line, useful when script run with linux nohup command
def print_flush(s):
    print s
    sys.stdout.flush()
#

# program execution starts here
print('Script to aggregate spamcount/unqips for *ALL* ASNs per year, and store to MySQL.')

# parse command line arguments - this scripts gets the year to process from the command line
if len(sys.argv) != 2:
    print '*** fatal error: missing argument(s). required arguemnts:  year'
    sys.exit()
glbl_yr = int(sys.argv[1])
print 'Running for year %d.\n' % glbl_yr

stt = time.time()  # variable to track script run-time
K UNROUTED = -1007 # dummy value for IPs from unknown ASNs
filecount = 0
asnlist = dict()       # holds pairs of asn:asndata

# main program loops - these outer loops run once for each file
for mo in range (1,13):
      print flush('Starting processing of month  %d...' % mo)
    for dy in range (1,32):
        # check if this combination is a valid day
        try:    today = date(glbl_yr,mo,dy)
        except: continue

        # try loading nearest geoipdb to current day
        tmp_geoipdb = 'geoip/GeoIP-106_%4d%02d%02d.dat' % (glbl_yr, mo, dy)
        if os.path.isfile(tmp geoipdb):
            gi = None
            gi = GeoIP.open(tmp_geoipdb, GeoIP.GEOIP_MEMORY_CACHE)
            current_gi = tmp_geoipdb[6:24]

        # open this day's file
        fname = "spam.daily/%4d/daily.%4d%02d%02d" % (glbl_yr, glbl_yr, mo, dy)
        try:
            fd = open(fname, "r")
            print_flush( '   .. %s  (w/ %s)  @t:%.1f' %(fd.name,current_gi.lower(),time.time()-stt) )
        except IOError:
            continue

        # get 'avg' figure for that day - the multiplier for spam_volume
        while fd.next() != "Log file retry statistics\n":
            pass
        fd.next()
        s = fd.next().split(',')[2].strip()
        avg = float ( s[s.find(' '):] )
        if avg < 1 or avg > 15:
            print '*** strange avg: ' + str(avg)
```

```
            avg = 4  # use 4 (in case of unusual values


        # now, go back to start of file to start processing records
        fd.seek(0)

        # the log file consists of two very similar sections. skip lines until section start
        for section in ['injection', 'logged']:
            section_start = "Overall %s points:\n" % section
            while fd.next() != section_start:
                pass
            fd.next()

             # this loop executes for each line until section finishes
             for s in fd:
                if s == '\n': break

                # parse line - the format is [ip count asn]
                ip = s[0:15].rstrip()
                count = int(s[16:21]) if section[0] == 'i' else int(s[16:25])
                pos = 33 if section[0] == 'i' else 42
                asn = int(s[pos:pos+7]) if s[pos].isdigit() else K UNROUTED

                cc = gi.country_code_by_addr(ip) or '--'        # perform GeoIP lookup!

                # get the ASNData object for this asn, or creates one if necessary
                # then add this record to it
                asnlist.setdefault(asn, AsnData(asn)).add_item(ip, cc, count*avg)
            #
        #
        fd.close()
        filecount += 1
    #
#

write_todb(asnlist)

print '####'
print 'files proccessed: %d\ntotal time: %.3f' % (filecount, time.time() - stt)
```

Conclusion

Due to the large quantity of raw data, a process of data preparation was necessary to summarize the data into a form suitable for statistical software packages. For this reason, a processing infrastructure was built during the course of the project (funded by the broader research this project was a part of). Other projects with similar processing needs can make use of concepts presented here.

One example application, mentioned by one of the supervisors of the author, was a project involving a major Dutch supermarket chain. The project required quantitative analysis of the logs produced by the cash registers. There, similar to our project, data on each individual item purchased is not useful for statistical analysis (and neither is it possible), but rather aggregations on various levels are needed (e.g., at the level of product categories, or the totals for different months or different neighborhoods).

4.2 Data Triangulation

4.2.1 Overview

It is important to know if our dataset is an 'indicative sample' of worldwide spam activity, before we can base generalized conclusions on it. For this reason, we need to compare our data with the public spam reports published by the industry. One of the (only) industry groups that have accurate data on spam statistics are the commercial security providers. For certain reasons, these companies refrain from making their data publicly available or even producing in-depth (non-marketing) analysis based on them.[1] The major security providers do however publish excerpts of their data, in the form of regular reports on malware and spam trends. It was decided to go through these publicly accessible reports (listed in Table 18), and find parts that could be *triangulated* with our data. Due to a variety of reasons, most of the information provided in these reports is not useful for triangulation purposes, leaving us with the following two possible comparisons:[2]

- Graphs of the global spam trends
- List of top spam emitting countries

Table 18 – List of the major publicly available security reports

Security Firm	Report Name
Cisco	Cisco annual security report
Google (Postini)	Annual Google communications intelligence report
IBM X-Force	IBM Internet security systems: X-Force trend & risk report
IronPort (part of Cisco)	Internet security trends Internet malware trends
Kaspersky Lab	Statistics – Kaspersky security bulletin
McAfee	McAfee research report - spam report
MessageLabs (part of Symantec)	MessageLabs intelligence: annual security report
Microsoft	Microsoft security intelligence report
Panda Security	PandaLabs Annual Report
Sophos	Security threat report
Symantec	Symantec Internet security threat report
Trend Micro	Trend Micro annual threat roundup and forecast

4.2.2 Comparison of spam trend graphs

In Figure 34 we have plotted next to each other the spam trends for 2007, based on our data, and from IronPort (2008b). By simply comparing the plots it can be seen that trends match quite nicely. Figure 35 presents the same plots for 2008. For this year, the match is not as good. Luckily, the difference in 2008 might not be problematic. In a workshop held in September 2009 with security experts to discuss these findings (and those of the broader OECD project it is part of), experts pointed out that this difference might be due to the rise of *'snowshoe spamming'* in

[1] The most probable reason is that publishing this information would outrage some of their worse performing customers. Another reason that comes to mind is to protect the location of their spam-traps and honey-pots.

[2] Some reasons why the information provided is not useful for triangulation include: i) reporting spam as a percentage of email rather than absolute numbers, making it non-comparable since we don't have any numbers of email growth; ii) reporting spam type (image, text, etc), message size, or intent (phishing, sales, etc), which we cannot compare since we don't have actual spam messages; iii) and finally, reporting case studies of particular threats, rather than global statistics, which makes sense for marketing purposes, but is not useful for us.

2008. Snowshoe spam is sent from static IP addresses belonging to a narrow range of IPs, and according to Spamhaus (2009), it rose in 2008, to be second only to botnet spam.

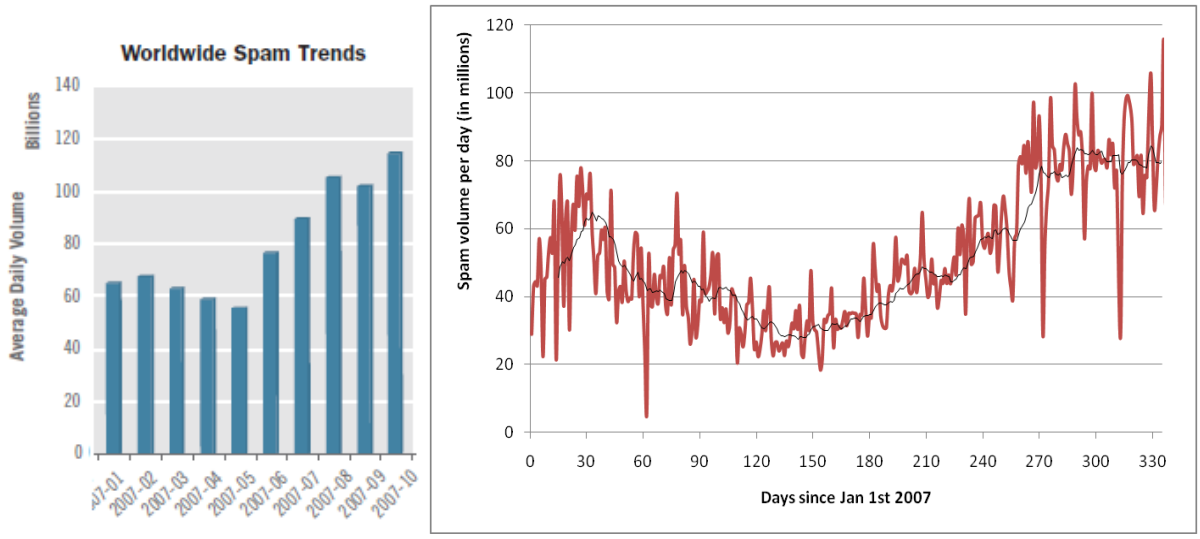

Figure 34 - Side by side comparison of worldwide spam trends for 2007 (left: IronPort (2008b), right: our data)

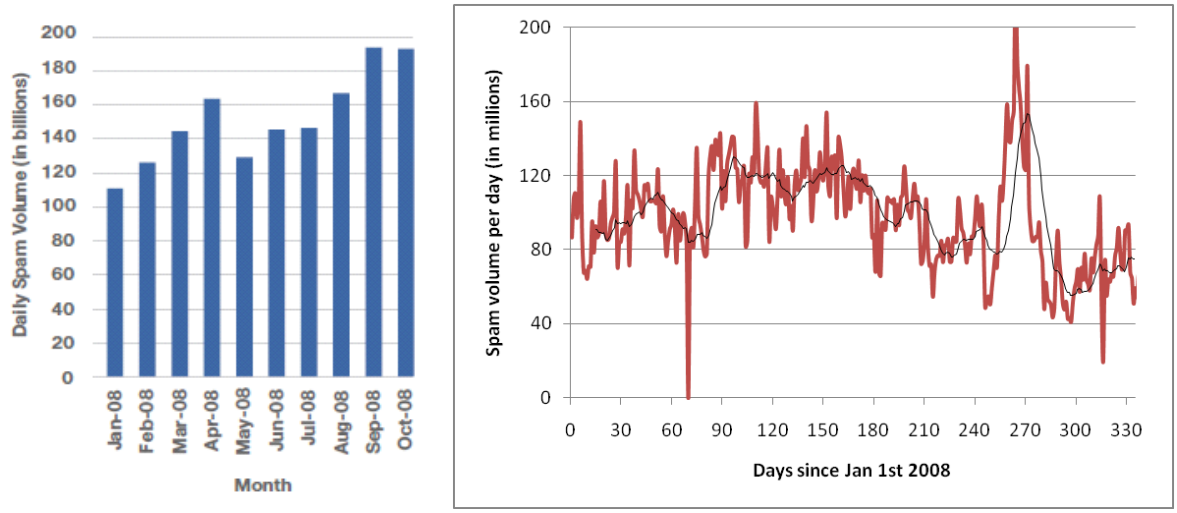

Figure 35 - Side by side comparison of worldwide spam trends for 2008 (left: Cisco (2008), right: our data)

Due to the use of static IPs, snowshoe spam has a higher chance of passing through spam filters. The technique evades IP reputation services in this way: back-end spam servers 'tunnel' their spam through the egress (static) IP addresses. Should the egress IP addresses become blacklisted, spammers shift to another range of addresses. The back-end servers remain hidden and undetected.[1] As this technique is costlier than *botnet spam* (due to the price paid for the IP addresses), snowshoe spammers employ up to date mailing lists, and avoid sending to stale email addresses or small domains – which is exactly the case of our spam-trap. This means that our spam-trap won't pickup snowshoe spam, and hence the spam trends it shows will not be indicative of worldwide spam trends. It however picks up *'botnet spam'* (and most probably remains indicative of that, could we test it). We are actually interested in botnet spam, so this difference is not problematic.

[1] These IP ranges are quite diverse. Due to the backend servers remaining undetected, these spammers usually provide the 'customer of customer' excuse to ISPs (Spamhaus 2009).

4.2.3 Comparison of top spam sending countries

The second triangulation attempt makes use of formal statistical tests. The basic idea is to see how well the list of top spam sending countries published in industry reports, matches the list generated from our data. Due to the limited number of observations, we need to use *non-parametric tests of association*. According to Siegel and Castellan (1988), the following statistical tests, among others, are suitable for this scenario:

- Spearman rank-order correlation coefficient (*Spearman's rho*)
- Kendall rank-order correlation coefficient (*Kendall's tau*)
- Kendall coefficient of concordance (*Kendall's W*) - also known as ranking N items by K judges.

Figure 36 displays the top spam sending countries in 2007 from two sources, KasperskyLab (2008) and IBM (2008). Table 19 lists the same data according to Sophos (2007), and our own dataset.

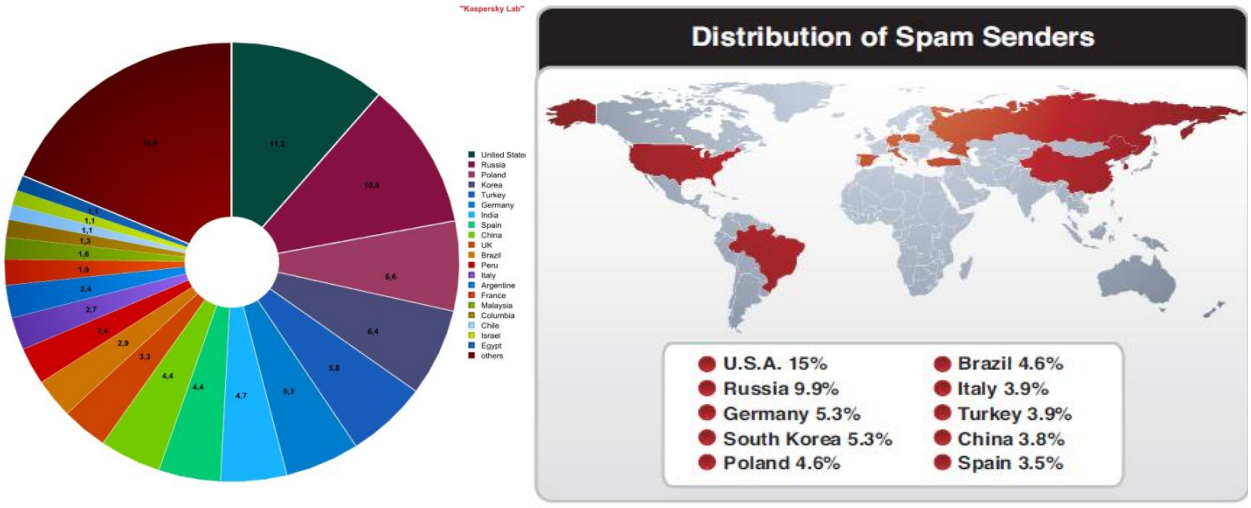

Figure 36 - Top spam emitting countries in 2007 - left: KasperskyLab (2008) , right: IBM (2008)

Table 19 - Top spam emitting countries in 2007 – Sophos (2007) , and our data

Sophos		Our data	
United States	22.5%	United States	17.5%
South Korea	6.5%	Russia	5.9%
China (incl HK)	6.0%	China	5.0%
Poland	4.9%	Poland	5.0%
Russia	4.7%	Brazil	4.7%
Brazil	3.8%	Germany	4.7%
France	3.5%	South Korea	4.6%
Germany	3.5%	France	3.7%
Turkey	3.1%	Italy	3.4%
Spain	2.7%	Turkey	2.8%
Italy	2.7%	UK	2.7%
India	2.6%	Spain	2.7%
		India	2.5%

The results of the statistical tests are presented below. As can be seen, at the 0.05 significance level, our data is associated with X-Force (ρ = 0.565) and Sophos (ρ = 0.719), but not with Kaspersky. Based on Kendall's W, the similarity between the four rankings is *not* rejected.

Correlations

			XForce	our_data	Kaspersky	Sophos
Kendall's tau_b	XForce	Correlation Coefficient	1.000	.565*	.582*	.489
		Sig. (2-tailed)	.	.029	.023	.056
		N	10	10	10	10
	our_data	Correlation Coefficient	.565*	1.000	.395	.719**
		Sig. (2-tailed)	.029	.	.065	.001
		N	10	13	13	12
	Kaspersky	Correlation Coefficient	.582*	.395	1.000	.388
		Sig. (2-tailed)	.023	.065	.	.084
		N	10	13	13	12
	Sophos	Correlation Coefficient	.489	.719**	.388	1.000
		Sig. (2-tailed)	.056	.001	.084	.
		N	10	12	12	12
Spearman's rho	XForce	Correlation Coefficient	1.000	.674*	.718*	.592
		Sig. (2-tailed)	.	.033	.019	.071
		N	10	10	10	10
	our_data	Correlation Coefficient	.674*	1.000	.528	.850**
		Sig. (2-tailed)	.033	.	.064	.000
		N	10	13	13	12
	Kaspersky	Correlation Coefficient	.718*	.528	1.000	.550
		Sig. (2-tailed)	.019	.064	.	.064
		N	10	13	13	12
	Sophos	Correlation Coefficient	.592	.850**	.550	1.000
		Sig. (2-tailed)	.071	.000	.064	.
		N	10	12	12	12

*. Correlation is significant at the 0.05 level (2-tailed).

**. Correlation is significant at the 0.01 level (2-tailed).

Kendall's W Test

Ranks

	Mean Rank
XForce	2.55
Our data	2.35
Kaspersky	2.80
Sophos	2.30

Test Statistics

N	10
Kendall's W[a]	.032
Chi-Square	.959
df	3
Asymp. Sig.	.811

Table 20 lists the spam emitting countries in 2008 from three industry reports and our data. The results of the statistical tests are presented beneath the table. Based on Spearman's rho, and at the 0.05 significance level, our data is associated with all the other three sources – IronPort (ρ = 0.620), Sophos (ρ = 0.935), and X-Force (ρ = 0.939). Based on Kendall's tau, our data is significantly associated with Sophos and X-Force, but not IronPort. Based on Kendall's W, the similarity between the four rankings is not rejected.

Table 20 - Top spam emitting countries in 2008, as listed in IronPort(2008b), Sophos (2008), X-Force / IBM (2009), and our own dataset

IronPort		*Our data*		*Sophos*		*X-Force*	
USA	17.2%	United States	13.4%	USA	17.5%	Russia	12.0%
Turkey	9.2%	Russia	9.0%	Russia	7.8%	USA	9.6%
Russia	8.0%	Brazil	6.4%	Turkey	6.9%	Turkey	7.8%
Canada	4.7%	Turkey	4.5%	China (inc HK)	6.0%	Brazil	5.6%
Brazil	4.1%	China	4.1%	Brazil	4.4%	China	4.4%
India	3.5%	Italy	4.1%	South Korea	3.7%	South Korea	4.0%
Poland	3.4%	South Korea	3.9%	Italy	3.3%	UK	3.3%
Korea	3.3%	Poland	3.4%	UK	3.1%	Spain	3.2%
Germany	2.9%	United Kingdom	3.2%	Poland	3.0%	Poland	3.2%
UK	2.9%	Spain	3.1%	India	2.9%	Germany	3.2%
Thailand	2.8%	Germany	2.9%	Spain	2.8%		
Spain	2.8%	Argentina	2.8%	Germany	2.7%		
Italy	2.4%	France	2.8%				
Argentina	2.1%	Colombia	2.6%				
Columbia	2.1%	India	2.4%				
France	2.0%	Romania	1.9%				

Correlations

			IronPort	Our_DS	Sophos	XForce
Kendall's tau_b	IronPort	Correlation Coefficient	1.000	.440	.550*	.706*
		Sig. (2-tailed)	.	.061	.019	.010
		N	13	11	11	9
	Our_DS	Correlation Coefficient	.440	1.000	.809**	.828**
		Sig. (2-tailed)	.061	.	.000	.001
		N	11	12	12	10
	Sophos	Correlation Coefficient	.550*	.809**	1.000	.874**
		Sig. (2-tailed)	.019	.000	.	.001
		N	11	12	12	10
	XForce	Correlation Coefficient	.706*	.828**	.874**	1.000
		Sig. (2-tailed)	.010	.001	.001	.
		N	9	10	10	10
Spearman's rho	IronPort	Correlation Coefficient	1.000	.620*	.729*	.834**
		Sig. (2-tailed)	.	.042	.011	.005
		N	13	11	11	9
	Our_ds	Correlation Coefficient	.620*	1.000	.935**	.939**
		Sig. (2-tailed)	.042	.	.000	.000
		N	11	12	12	10
	Sophos	Correlation Coefficient	.729*	.935**	1.000	.963**
		Sig. (2-tailed)	.011	.000	.	.000
		N	11	12	12	10
	XForce	Correlation Coefficient	.834**	.939**	.963**	1.000
		Sig. (2-tailed)	.005	.000	.000	.
		N	9	10	10	10

*. Correlation is significant at the 0.05 level (2-tailed).

**. Correlation is significant at the 0.01 level (2-tailed).

Kendall's W Test

Ranks

	Mean Rank
Our_DS	2.67
Sophos	1.94
XForce	3.22
IronPort	2.17

Test Statistics

N	9
Kendall's W[a]	.201
Chi-Square	5.414
Df	3
Asymp. Sig.	.144

Conclusion

A prerequisite for the generalizability of our final results is to validate that our spam data is a representative sample of worldwide spam trends, e.g., by triangulating it with other available sources. Unfortunately, due to the scarcity of public data, only a limited number of triangulation tests are possible.

Comparing the spam trend graphs shows a very good match for 2007, and a not as good, but acceptable situation for 2008. Comparison of the top spam senders, our list associates well with some of the industry reports, but not with some others. Fortunately, we happen to be in the middle of the pack (as the reports do not agree among themselves either).

Putting these together, we believe that our data is a valid representative sample. [1]

[1] Of course, practicality in research tells us that even if the triangulation results were not good, we would have had no choice but to stick with the data we have (and simply accept the difference as a limitation)!

4.3 Final Steps

4.3.1 Selecting the ISPs to analyze

We are almost ready to start the data analysis. An important remaining step is the choice of ISPs to include in the final dataset. Out of whole set of retail ISPs in the TeleGeography database, it was decided to focus only on ISPs that operate in the so called *extended OECD countries*. The OECD, or Organization for Economic Cooperation and Development, is an international organization of countries that are "committed to democracy and the market economy" (OECD 2009). Most of these countries are regarded as high income and developed countries. The OECD consists of 30 members, and in addition, has five 'accession candidates' and another five 'enhanced engagement' partners. These countries, together with the number of ISPs from each in our dataset, are listed in Table 21.

The reason for focusing on a limited set of countries is the laborious work involved in mapping ASNs to operators, and in gathering values for the independent variables. The extended OECD countries are a good subset for two reasons: it includes a list of states for whom reliable statistics is more readily found, and these states includes all the important state, from both an economic perspective, and an Internet usage perspective.

Table 21 - 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**

4.3.2 Mapping ASNs to operators

The process of mapping ASNs to operators has been executed as follows:

1. The amount of spam sources and spam volume was generated per ASN/CC, worldwide, for all the years.
2. For each of the selected countries, the list of ASNs was sorted by the percentage of spam sources / volume seen from that particular ASN. All ASNs that were above the 0.5% threshold were noted of. (The number of ASNs chosen differs per countries – e.g., 20 for the Netherlands, and 44 for the U.S.).
3. With the help of Mr Menno Nederveen (a student assistant at the section), the name of each of the noted ASNs was searched for on Google, Wikipedia, ISP sites, etc, to see which of the TeleGeography operators it matches.

In some cases, the mapping was pretty straightforward. Other times, such as cases of mergers and acquisitions, the ASN was mapped to its current owner. For a limited number of operators, no corresponding ASN was found, which seems unlikely (this could be due to the ASN falling under the 0.5% threshold). The whole process has been carried out rather carefully, but the possibility of mistakes exists. We do believe that the margin of error is not high. In any case, this manual mapping process and the mistakes that come with it have to be accepted as a limitation of this (and similar) research.

4.3.3 Final dataset

The final Stata dataset has 741 observations. The breakup of these observations based on year is presented in Table 22. The list of variables and the count of observations for each is presented in Table 23.

Table 22 - Observations in dataset

Year	# of observations
2005	170
2006	182
2007	195
2008	194

Table 23 - Variables in dataset

Variable	Obs
opcode	-
op_name	-
year	741
country	-
unq_srcs	741
spam_msgs	741
src_persub	741
spam_persub	741
total_sub	741
market_share	709
rev_persub	194
srv_cable	665
cyberconv_mem	741
lap_mem	741
educ_ix	547
int_bpp	386
piracy_rate	740

CHAPTER 5 – DATA ANALYSIS

5.1 OVERVIEW

In the previous chapters, we developed a set of hypotheses, which when answered will shed some light on how ISPs are mitigating botnets; in the broader picture, this will eventually lead to designing mechanisms for fighting off this growing phenomenon. We also explained the process through which we build a dataset that will aid us in testing our hypotheses; a dataset which is very rich in certain variables, and somewhat limited in others (forcing us to adopt many proxies). Despite its limitations however, we believe that this data, if analyzed thoroughly, has the potential to tell us quite a lot – most important of all, to empirically test some of the assumptions being made regarding the botnet phenomenon in the literature. In order to do this, we will employ a variety of statistical instruments to test the hypotheses, as shall be explained.

5.1.1 STATISTICAL INSTRUMENTS

We have nine hypotheses to test (relisted in Table 24); all of them can be tested on their own using a variety of statistical techniques. The general procedure that we will follow is presented in Figure 37.

The first part of our analysis (section 5.2) will use statistical tests performed individually for each hypothesis. For H1 and H2, which only use the dependent variable, descriptive statistics will be used. For H3 – H9, which include independent variables, techniques involving comparison of means, and those involving measures of association will be deployed. Mode

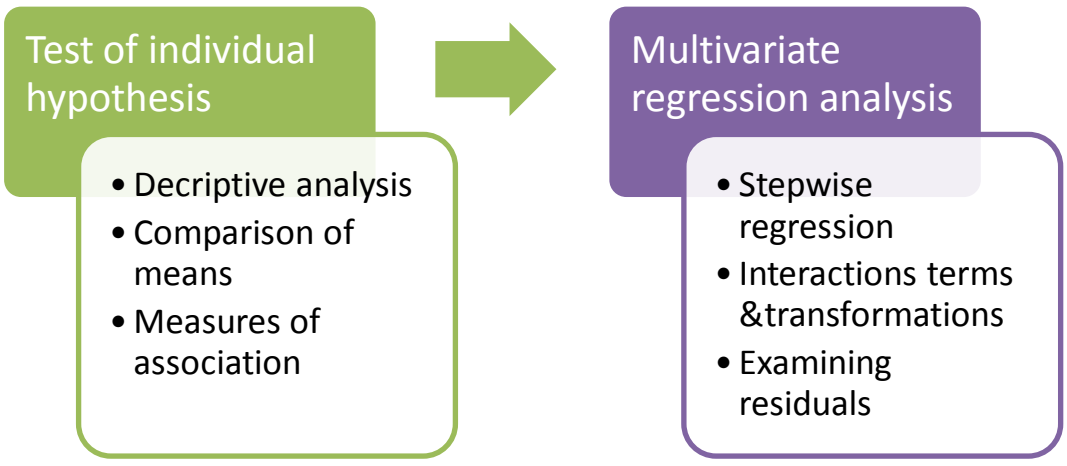

Figure 37 - Data analysis procedure[1]

[1] Another form of analysis was also attempted: transforming the variables into categorical scales, and thereafter using techniques such as contingency tables, partial tables, etc. This exercise did not produce fruitful results and was not continued.

The second part (section 5.3) will undertake building a full model that simultaneously includes all the independent variables. The reason for including this part is that although we will have already tested our hypothesis in the first part and identified various factors that influence the level of botnet activity, we will not know the relative importance of these factors when compared together, or when controlled for each other. More importantly, we wish to know what percentage of the variance between ISPs these factors together explain (i.e., what percentage of factors have been identified and what percentage is missing). For these reasons, we will build a full regression, model and thereafter assess it.

Table 24 – List of hypotheses to test

#	Hypothesis
H1	A majority of the world's malicious traffic originates from "autonomous systems" run and controlled by a limited number of ISPs, as compared to originating from ASes controlled by other types of organizations.
H2	ISPs significantly differ in terms of their performance in mitigating botnets.
H3	Larger ISPs perform worse in terms of security (i.e., have a lower security performance)
H4	ISPs with higher average revenue per user have higher security performance
H5	Cable providers have higher security performance than DSL providers
H6	ISPs in countries that have endorsed international agreements against cyber-crime (e.g., the LAP or the cyber-crime convention) have higher security performance
H7	ISPs in countries with higher piracy rates have lower security performance
H8	ISPs in countries with higher average bandwidth rates have lower security performance
H9	ISPs in countries with a higher educational index have higher security performance

5.1.2 POOLED DATA VERSUS FOCUSING ON A SINGLE YEAR

As explained in the previous chapters, our unit of analysis are Internet service providers and network operators active in the 40 enlarged OECD[1] countries. Our observations span the years 2005 to 2008, (as listed in Table 22 of the previous chapter, with the total adding up to 741.

It should be noted that although in most cases having more years in a dataset produces more powerful generalizations, in our case the opposite might actually be true, given the dynamic nature of our phenomenon. Put another way, a pattern that holds in 2005 might not hold in 2008; thus combining all the years together would actually obscure the detection of such patterns, instead of making them bolder. For this reason, we will use two versions of the datasets in the analysis: one pooled, containing all the observations, and the other with only 2007 data. The year 2007 was chosen because at the time of this writing, some of the global indicators for 2008 have yet to be published.

[1] The enlarged OECD consists of the 30 member states, 5 accession candidates, and 5 enhanced engagement partners.

5.2 Individual Tests of Hypothesis

This section consists of sub-sections that each test one of the hypotheses. Please note that since we need to perform each test for two dependent variables (*spam sources* and *spam volume* - as was decided the methodology), and for two datasets, we end up with four sets of statistical results for each sub-section. Thus, we include at the end of each sub-section a heading that compares these results and make a final judgement on the acceptance or rejection of the hypothesis. Please be advised that this section is rather verbose, with detailed Stata outputs and many figures (histograms, box-plots, scatter-plots). Should the reader want, she can skip directly to section 5.2.10 which presents the summary of the findings in one table.

5.2.1 Hypothesis 1: ISPs are central

Our first hypothesis, which is basically an answer to the question *"Are ISPs central (to the botnet problem)?"*, is stated as follows:

> ***A majority of the world's malicious traffic originates from "autonomous systems" run and controlled by a limited number of ISPs, as compared to originating from ASes controlled by other types of organizations.***

This hypothesis can be answered using descriptive statistics, based on the broader dataset. We can simply show that in the OECD countries, more than 80% of the infected sources sending out spam lie within the approximately 400 ASNs operated by the 200 bigger ISPs. This figure is astonishing, considering that there are a total of 30,000+ active ASNs.[1] The actual figure differs in each country, going all the way up to 98% for countries such as Turkey and France. Figure 38 shows the percentages for each of the countries in 2007. The blue bars (left-hand bars) represent the percentage of infected sources located in this subset of ASNs, and the red bars show the proportion of spam these ASNs are responsible for. The average of the blue bars is 81% with a standard deviation of 12%.

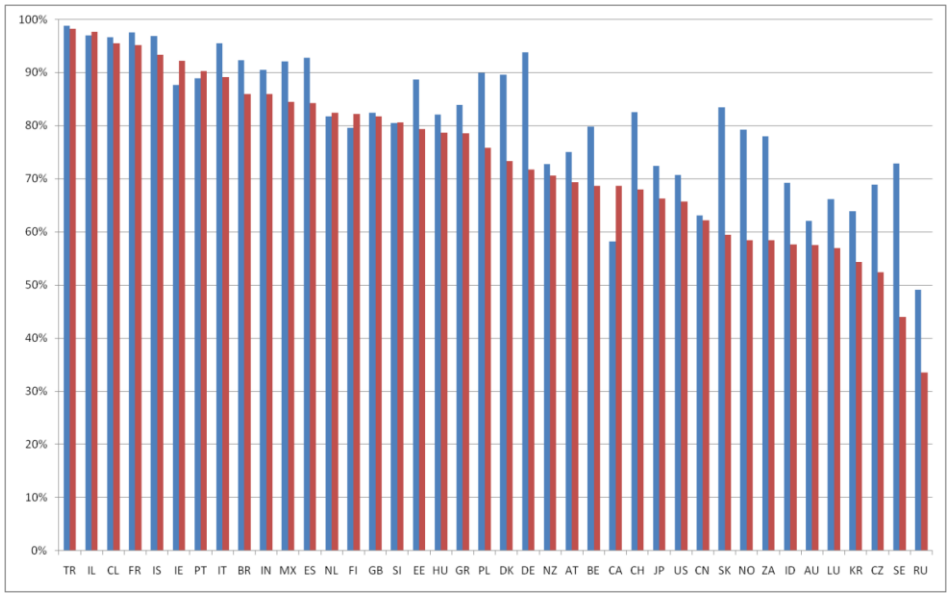

Figure 38 - Percentage of spam emmitted by the 200 major OECD ISPs, and percentage of infected sources located in each (by country, 2007)

[1] Yet another example of the 80/20 principle!

Figure 39 shows this percentage of infected sources for the years 2005 till 2008 (totals within the OECD area), and Table 25 shows the breakup of these numbers (all based on our dataset).

Table 25 - Number of infected sources, and spam messages emmitted annually, worldwide / OECD countries / top 200 ISPs

	Unique sources (global total)	Unique sources (OECD total)	Unique sources (200 ISPs)	Top ISPs to global ratio	Top ISPs to OECD ratio
2005	22,381,514	19,891,261	15,806,148	71%	79%
2006	44,234,200	38,504,784	31,074,776	70%	81%
2007	73,209,230	62,284,870	50,890,908	70%	82%
2008	66,696,170	53,635,240	42,815,109	64%	80%
	Spam messages (global total)	**Spam messages (OECD total)**	**Spam messages (200 ISPs)**	**Top ISPs to global ratio**	**Top ISPs to OECD ratio**
2005	3,204,615,662	2,890,349,417	2,026,048,665	63%	70%
2006	4,860,608,032	4,200,006,085	3,207,378,700	66%	76%
2007	19,931,685,581	15,787,838,093	11,366,137,646	57%	72%
2008	34,922,695,193	26,882,390,272	17,779,945,673	51%	66%

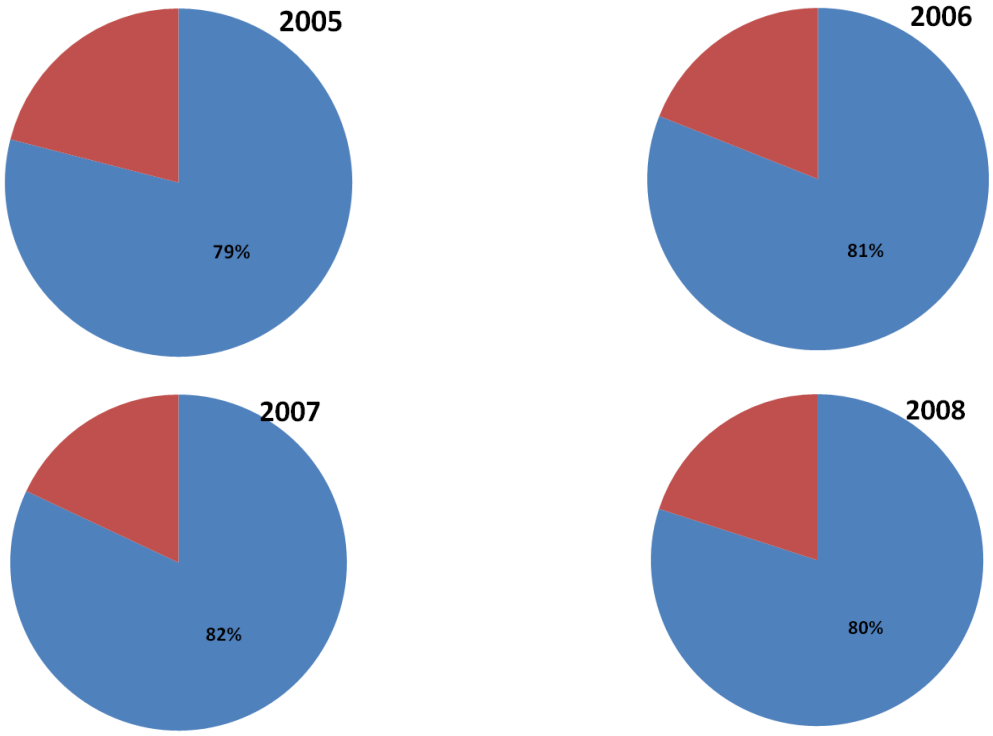

Figure 39 - Percentage of infected sources (in OECD countries) that are located in the 200 (predominantly) retail ISPs, by year

Finding

As we saw from our data, 80% of infected sources are located in a ASNs owned by 200 ISPs. Put another way, these 200 organizations are in a position to control 80% of the OECD's botnet activity; this is truly amazing, and shows that ISPs are indeed gatekeepers of the botnet problem. This is good news for the policy makers: bodies who would like to govern the botnet problem need only to negotiate with a limited number of actors (instead of for instance having to focus on millions of end users.)

5.2.2 HYPOTHESIS 2: ISPS DIFFER SIGNIFICANTLY

ISPs significantly differ in terms of their performance in mitigating botnets.

Testing this hypothesis requires examining the dispersion (variance) of the dependant variable in our sample. We proposed several dependant variables in the previous chapters, all of them a measure of botnet activity in an ISP network, which included *src_persub, unq_srcs, spam_persub,* and *spam_msgs* (Lower values of these variables indicate a higher performance in mitigating botnets). We will use descriptive statistics to examine the distribution and dispersion of these variables, and see whether there is a big variation between them or not.

2007 DATA

Unique sources, per subscriber (src_persub)

This is our main candidate for measuring botnet activity. It takes into account "unique spam emitting source" being approximately equivalent to an infected machine, and corrects this figure for the number of subscribers on the ISP network. Roughly speaking, its value indicates the percentage of compromised machines in an ISP in a certain period.[1] The summary statistics for this variable is as follows:

```
    variable |         N      mean        sd       min       max        cv
-------------+--------------------------------------------------------------
  src_persub |       196  .2300474  .2362841     .0002    1.1329   1.02711
-------------------------------------------------------------------------------
```

A common measure of dispersion is the ***coefficient of variation***, which is basically the ratio of the standard deviation to the mean**.** For our sample, this value is 1.02. Distributions with a CV of ≥ 1 are considered to have a high-variance. The histogram is displayed in Figure 40. An exponential distribution can be observed.

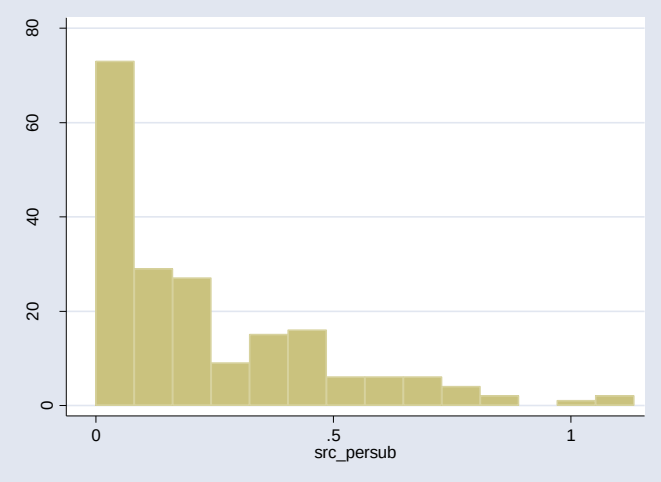

Figure 40 - Histogram of *src_persub* (2007)

[1] We say 'roughly speaking', due to the various technical limitations this metric has, such as dynamic IPs, NATs, etc (see Ch. 3)

Unique sources, absolute (unq_srcs)

We have previously argued the need to take into account the size of an ISP when comparing the number of bots residing within them. Despite those arguments, looking at the absolute number of bots (i.e., not relative botnet activity) has two benefits. If we plot the number of bots (unique sources) versus an ISP's size (subscribers), we will first of all, understand if size really matters (empirically), and second, by looking at the so called 'bandwidth', we can examine the variation in botnet activity in ISPs of the same size.

Figure 41 shows this scatter plot. We can already see that for certain ISPs with similar number of subscribers we see alternating levels of infected sources. Since the graph is condensed in the bottom left corner, a version with logged axes is also presented in the figure. Two things can be clearly seen from this figure. First is the fact that a positive relationship exists between total_sub and unq_srcs (i.e., larger ISPs tend to have more bots). Second is that for ISPs of similar size, there exists an order of magnitude of 2 difference (=100x)! [1] The key point is that the variation is quite large. Table 26 lists and compares these absolute numbers for several ISPs of similar size (cases have been selected to highlight extremes).

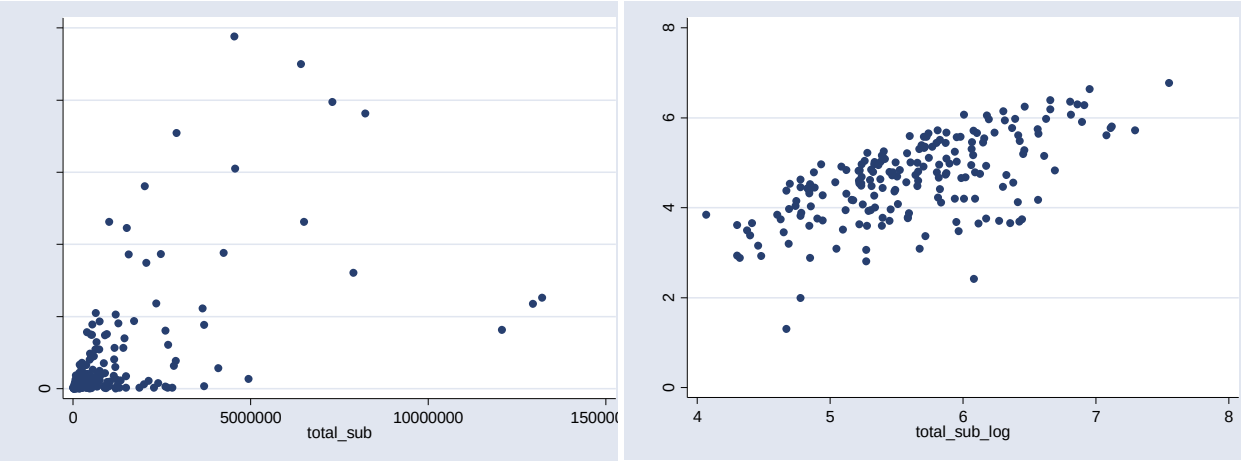

Figure 41 - Scatter plot of unq_srcs vs. total_sub (2007) – left: normal with outliers removed; right: double logged

Table 26 – Sample comparison of the number of infected sources in ISPs of similar size (Q4 2008)

Group	ISP	Subscribers	Unique Sources
Small (~500,000 subscribers)	CA03	495000	5417
	IL04	566000	96990
	RU06	500000	**164676**
	FI05	478000	**582**
Medium (~2.8 million subscribers)	BR06	2557800	**574597**
	DE01	3003000	152095
	DE11	2820000	2823
	IN01	3007415	340586
	JP02	2616000	**709**
	NL03	2655000	9525
	US18	2847000	40273
Large (~6-8 million subscribers)	FR02	8347618	**11289**
	TR01	5800000	**2003704**
	IT02	6754000	660208

[1] Putting these two points together, we can state that although the number of subscribers is one of the *main* driving factors for botnet activity in an ISP (as measured by the absolute number of bots). There still remain *considerable* differences between ISPs in terms of security performance (that can be rooted back to differences in security policies, and other factors that we are interested in uncovering).

Spam messages, per subscriber and absolute (spam_persub and spam_msgs)

We can run duplicate the same descriptive statistics that we just did for the number of spam messages. Let us start with *spam_persub*. The descriptive summary is as follows:

variable	N	mean	sd	min	max	cv
spam_persub	196	72.76714	100.8644	.33	1078.22	1.386126

The coefficient of variation is again ≥ 1, indicating a high variance in the sample. The histogram is shown in Figure 42, and the exponential distribution can again be observed.

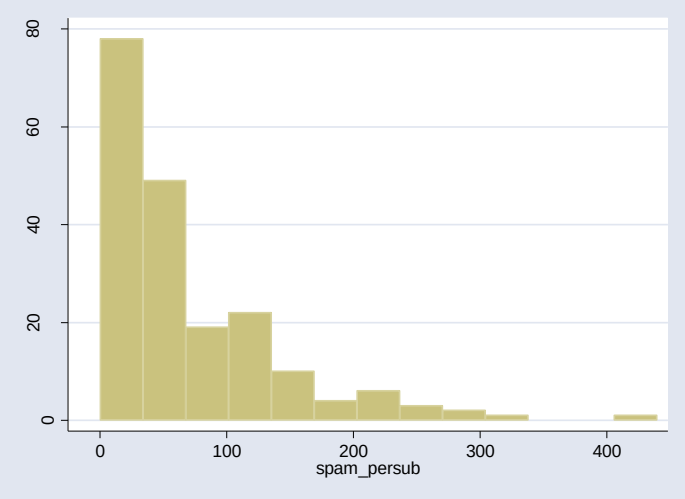

Figure 42 - Histogram of spam_persub (2007) - one outlier removed

For the absolute variable *spam_msgs*, we present the scatter plot versus the number of subscribers, and can observe a similar spread among ISPs of similar size, as presented in Figure 43. We see up to 3 orders of a magnitude difference here (=1000x), again pointing to a big variation between ISPs.

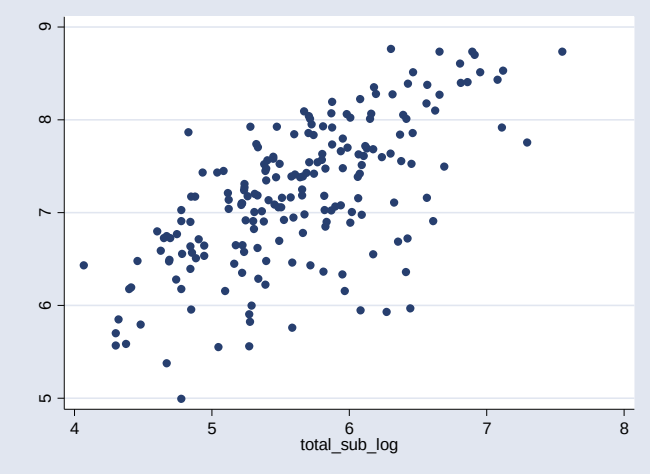

Figure 43 - Scatter plot of LOG of spam_msgs versus LOG of total_sub (2007)

POOLED DATA

Since the explanations for the analysis of the pooled data are similar to the 2007 data, we will simply present the results, and give explanations only were a difference exists.

Unique sources, per subscriber and absolute (src_persub and unq_srcs)

The summary statistics for this variable are presented below. As can be seen, the coefficient of variance is similarly ≥ 1, indicating a high variance. The histogram is displayed in Figure 44, and is likewise, an exponential distribution.

variable	N	mean	sd	min	max	cv
src_persub	741	.191359	.2107284	.0001	1.1329	1.10122

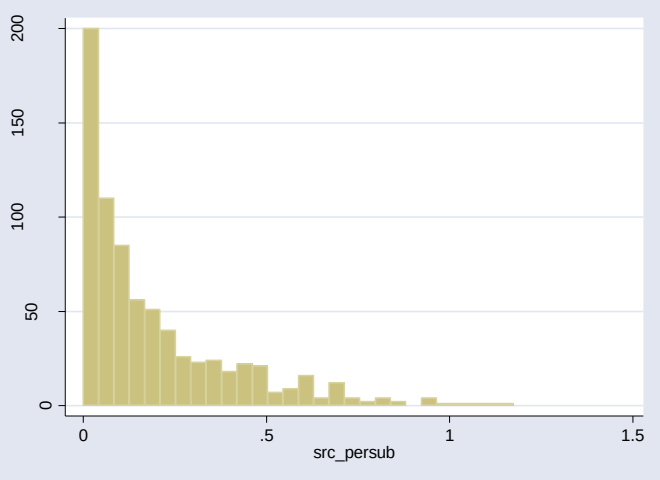

Figure 44 - Histogram of src_persub (2005-2008)

The scatter plot of the number of bots (unique sources) versus an ISP's size (subscribers) is presented in a log scale in Figure 45. The 'bandwidth' in the pooled data also shows an order of magnitude of 2 difference, between ISPs of similar size (x100). The variation is rather large.

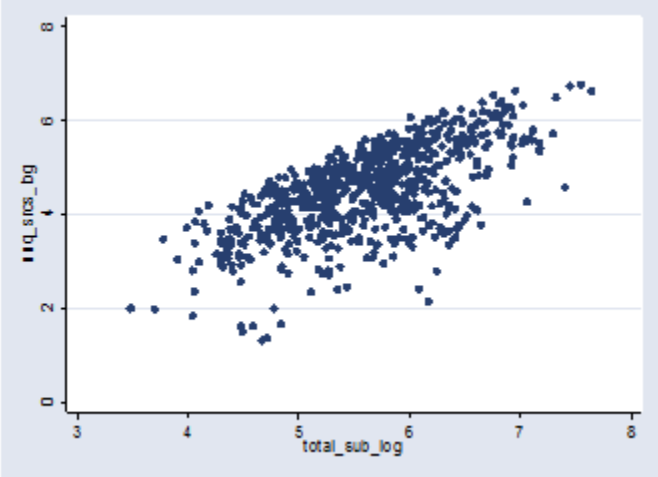

Figure 45 - Scatter plot of unq_srcs versus total_sub, logged (2005-2008)

Spam messages, per subscriber and absolute (spam_persub and spam_msgs)

The descriptive statistics for the variable *spam_persub* in the pooled data is presented below. The coefficient of variance is again ≥1, indicating a high variance.

variable	N	mean	sd	min	max	cv
spam_persub	741	55.57274	79.03818	.1697	830.8522	1.422247

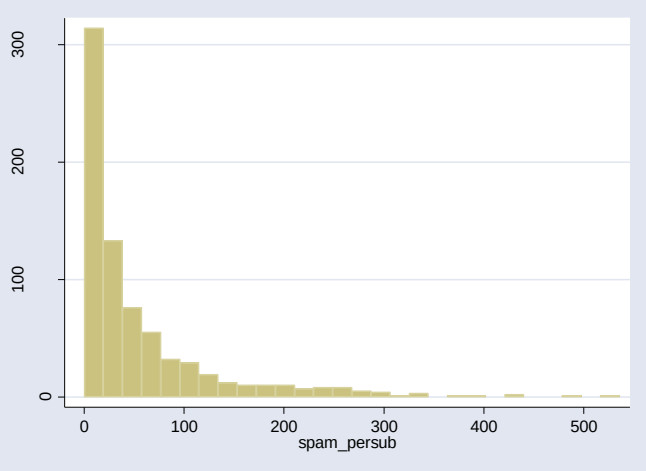

Figure 46 - Histogram of spam_persub (2005-2008) - outliers removed

For the absolute variable *spam_msgs*, we present its scatter plot versus the number of subscribers in Figure 47 (logged). As we have consistently seen so far, a wide spread among ISPs of similar size exists yet again, with up to 3 orders of a magnitude difference (=1000x).

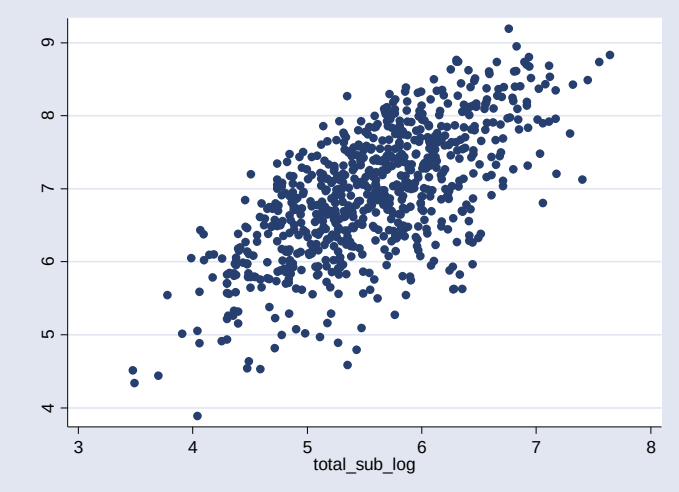

Figure 47 - Scatter plot of LOG of spam_msgs versus LOG of total_sub (2005-2008)

Finding

We can confidently accept the hypothesis that there is a big variation among ISPs in regards to security performance. This verdict is based on the coefficient of variance being ≥ 1 for src_persub and spam_persub, and the scatter plots of the absolute metrics versus size showing orders of magnitude spread, in both of our datasets.

5.2.3 HYPOTHESIS 3: EFFECTS OF ISP SIZE

Larger ISPs perform worse in terms of security (i.e., have a lower security performance).

Test strategy

Statistical tests to measure the *degree of association* between the following variables needs to be used:

- *src_persub* or *spam_persub* (botnet activity levels / security performance)
- *total_sub* or *market_share* (ISP size[1])

For this purpose, different tests can be used. If certain assumptions (namely, normality) are met, the ***Pearson correlation coefficient test*** can be used. If not, the non-parametric ***Spearman's rank correlation coefficient*** test can be used.[2] Additionally, scatter plots can give graphical sense to these associations.

2007 DATA

The descriptive statistics for the variables are as follows. The histograms are shown in Figure 40 and Figure 48. As it is observable from the Stata output, and the histograms, all four variables fail the test of normality.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
   total_sub |       196     1450039     3501650      11700   3.57e+07
market_share |       196    .1791832    .1977765       .001     1.0053

                  Skewness/Kurtosis tests for Normality
                                                 ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+-------------------------------------------------------
  src_persub |      0.000          0.004          35.68        0.0000
 spam_persub |      0.000          0.000              .        0.0000
   total_sub |      0.000          0.000              .        0.0000
market_share |      0.000          0.000          54.37        0.0000
```

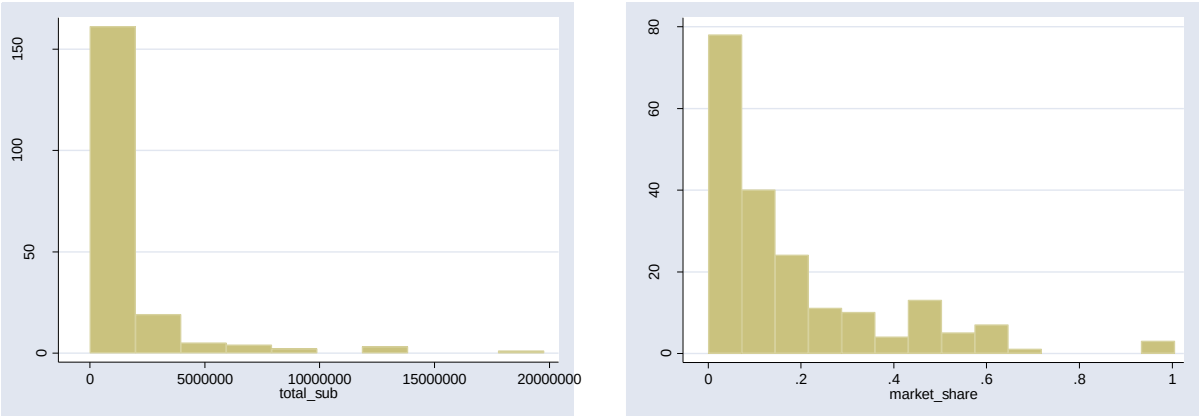

Figure 48 - Histogram of total_sub and market_share (2007)

[1] Please note that market_share is a very indirect proxy for size

[2] Another non-parametric measure of association, Kendall's rank correlation coefficient, yields similar results.

Test results

Due to the lack of normality, the non-parametric test to measure the degree of association between ISP size and botnet activity needs to be used. The null and alternate hypotheses for these statistical tests are presented below, and the actual test results in the box that follows:

H_0: botnet_activty and size_variable are independent (ρ=0)
H_α: botnet_activty and size_variable are positively correlated (ρ>0)

```
. spearman src_persub market_share
Number of obs =     196
Spearman's rho =       0.0377
Test of Ho: src_persub and market_share are independent
    Prob > |t| =       0.5995

. spearman spam_persub market_share
Number of obs =     196
Spearman's rho =      -0.0045
Test of Ho: spam_persub and market_share are independent
    Prob > |t| =       0.9496

. spearman src_persub total_sub
Number of obs =     196
Spearman's rho =      -0.1153
Test of Ho: src_persub and total_sub are independent
    Prob > |t| =       0.1075

. spearman spam_persub total_sub
Number of obs =     196
Spearman's rho =      -0.2606
Test of Ho: spam_persub and total_sub are independent
    Prob > |t| =       0.0002
```

As can be seen, taking market_share as the proxy for size, the result is statistically insignificant at the 0.05 level. If we take total_sub as the proxy for size, the result still remains insignificant for src_persub at the 0.05 level, but clearly stronger then with market share. For spam_persub, the result turns significant, but contrary to our prediction, with a negative rank correlation coefficient. That is, the smaller ISPs are actually doing worse in terms of security performance! In short: depending on the variables used, either ISP size and botnet activity are unrelated, or they are associated but in the reverse of our hypothesis. Figure 49 depicts the point that the biggest ISPs actually have lower botnet activity levels (compared in the relative 'per subscriber' metrics).

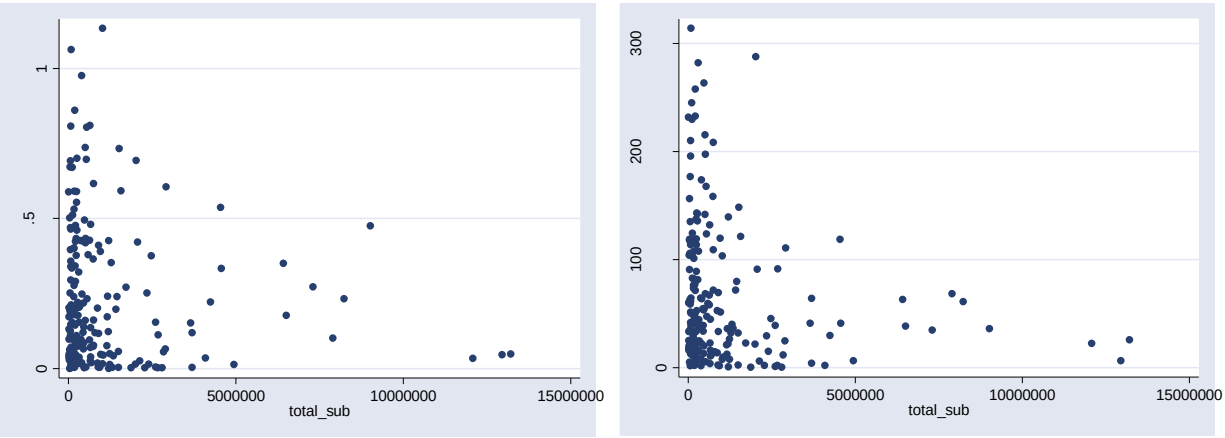

Figure 49 - Scatter plot of src_persub / spam_persub vs. total_sub (2007) – outliers removed

POOLED DATA

The descriptive statistics for the variables are as follows. As it is observable from the Stata output and the histograms, all four variables fail the test of normality. So again, a non-parametric test needs to be used.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
   total_sub |       741     1374231     3369101       3000   4.43e+07
market_share |       709    .1820523    .1988902      .0005     1.2358


                  Skewness/Kurtosis tests for Normality
                                                       ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)     Prob>chi2
-------------+-------------------------------------------------------------
  src_persub |      0.000         0.000               .        0.0000
 spam_persub |      0.000         0.000               .        0.0000
   total_sub |      0.000         0.000               .        0.0000
market_share |      0.000         0.000               .        0.0000
```

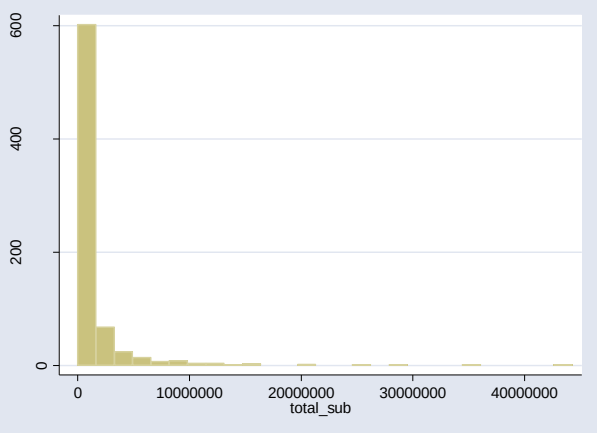

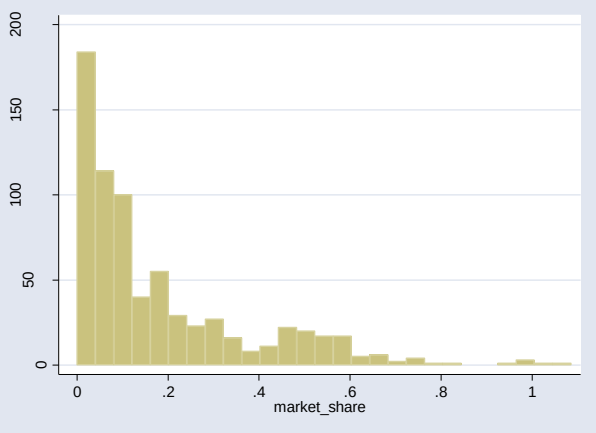

Figure 50 - Histogram of total_sub and market_share (2005-2008)

Test results

The null and alternate hypotheses for Spearman's rank correlation coefficient test are presented below, and the test results in the box that follows:

H_0: botnet_activty and size_variable are independent (ρ=0)
H_α: botnet_activty and size_variable are positively correlated (ρ>0)

```
. spearman src_persub total_sub
Number of obs =     741
Spearman's rho =      -0.1702
Test of Ho: src_persub and total_sub are independent
    Prob > |t| =       0.0000

. spearman spam_persub total_sub
Number of obs =     741
Spearman's rho =      -0.1816
Test of Ho: spam_persub and total_sub are independent
    Prob > |t| =       0.0000
```

```
. spearman src_persub market_share
Number of obs =     709
Spearman's rho =       0.0085
Test of Ho: src_persub and market_share are independent
    Prob > |t| =       0.8203

. spearman spam_persub market_share
Number of obs =     709
Spearman's rho =      -0.0505
Test of Ho: spam_persub and market_share are independent
    Prob > |t| =       0.1789
```

Again, taking market_share as the proxy for size, the results are statistically insignificant at the 0.05 level (and much lower levels). Interestingly enough, and similar to what we saw for the 2007 data, using total_sub as the proxy for size, we have a significant rank correlation, but in the opposite of the direction hypothesized. Figure 51 shows the scatter plots for these associations.

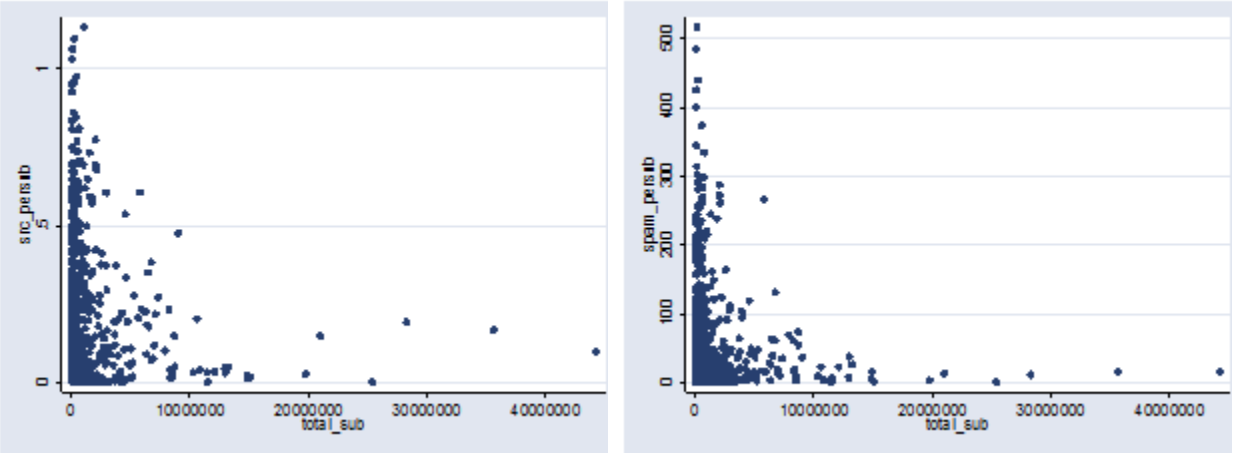

Figure 51 - Scatter plot of src_persub / spam_persub vs. total_sub (2005-2008) – outliers removed

Finding

Based on all the above findings, the third hypothesis is rejected. That is, contrary to what is believed and cited in the literature, we do not find empirical evidence to suggest that larger ISPs perform worse in regards to botnet mitigation.

What's more, when the number of subscribers of an ISP is used as the measure of size, we uncover an association in the reverse direction of our hypothesis. That is, the smaller ISPs have on average higher levels of botnet activity (relative to their number of subscribers), with $\rho \approx -0.17$.

The reasons for this could be that the larger ISPs typically use automated anti-spam/anti-botnet solutions that are more effective than manual approaches. It could also be that the weight of incentives such as reputation (of large brands) is more than the decrease in peer pressure noticed by the large operators. We will revisit these arguments in our conclusion chapter.

5.2.4 HYPOTHESIS 4: EFFECTS OF ARPU

ISPs with higher average revenue per user have higher security performance.

Test strategy

Statistical tests to measure the degree of association between the following variables needs to be used:

- *src_persub* or *spam_persub* (botnet activity levels / security performance)
- *rev_persub* (average revenue per user)

For this purpose, the non-parametric ***Spearman rank correlation coefficient test*** will be used (as we already know that the assumption of normality does not hold.) Scatter plots will also be drawn.

2007 DATA

The summary statistics, histograms, and test of normality for the independent variable *rev_persub* are as follows. The distribution for this variable is <u>not</u> normal. (Please note that the number of observations is about the third of the total, as revenue figures are not reported for many operators.)

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
  rev_persub |        59    3417.214    2711.386     527.76   14726.59

                  Skewness/Kurtosis tests for Normality
                                                 ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+--------------------------------------------------------
  rev_persub |      0.000         0.000          27.31        0.0000
```

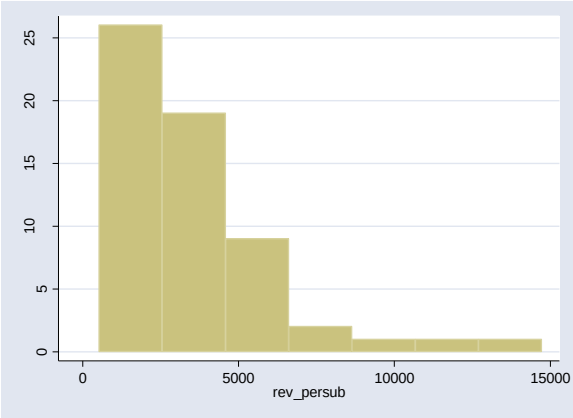

Figure 52 - Histogram of rev_persub (2007)

Test results

Due to the lack of normality, the non-parametric tests to measure the degree of association are used. The null and alternate hypotheses for these tests, and the actual test results, are as follows:

H_0: botnet_activty and rev_persub are independent (ρ=0)
H_α: botnet_activty and rev_persub are positively correlated (ρ>0)

```
. spearman src persub rev persub
Number of obs =      59
Spearman's rho =       0.1076
Test of Ho: src persub and rev persub are independent
    Prob > |t| =       0.4171

. spearman spam_persub rev_persub
Number of obs =      59
Spearman's rho =       0.1160
Test of Ho: spam_persub and rev_persub are independent
    Prob > |t| =       0.3817
```

As can be seen, the test results are statistically insignificant at the 0.05 level (and in addition, show a very low degree of association). So, the null hypothesis cannot be rejected, and we can assume that an ISP's average revenue per user and its botnet activity levels are unrelated. The scatter plots shown in Figure 53 illustrate this.

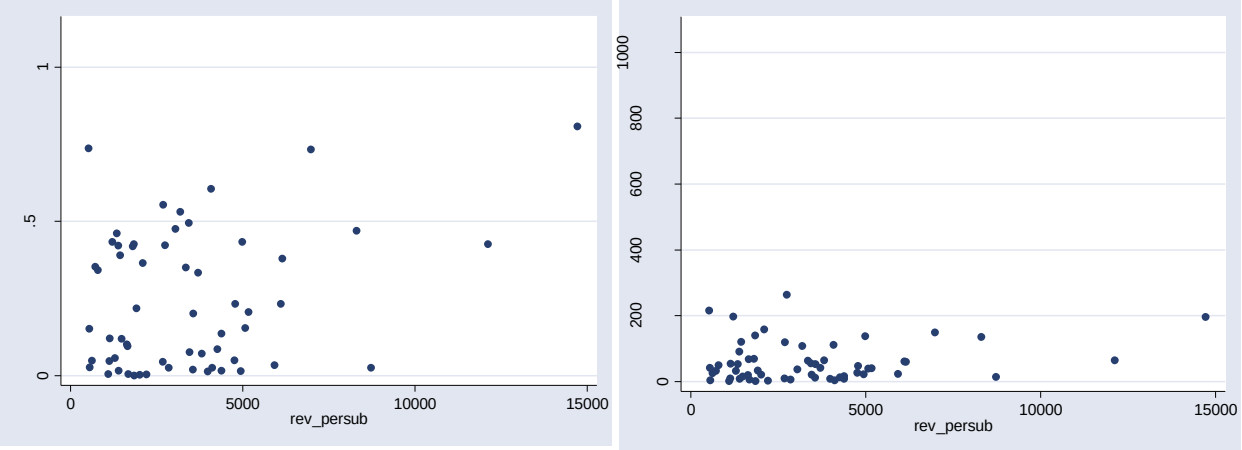

Figure 53 - Scatter plots. left: src_persub vs. rev_persub; right: spam_persub vs. rev_persub (2007)

POOLED DATA

The summary statistics, histograms, and test of normality for the independent variable *rev_persub* are as follows. The distribution for this variable is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
  rev_persub |       194    4305.685    5135.761   182.1285   42768.34


                  Skewness/Kurtosis tests for Normality
                                                 ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+--------------------------------------------------------
  rev_persub |      0.000         0.000           27.31       0.0000
```

Test results

Due to the lack of normality, the non-parametric tests to measure the degree of association are used. The null and alternate hypotheses for these tests, and the actual test results, are as follows:

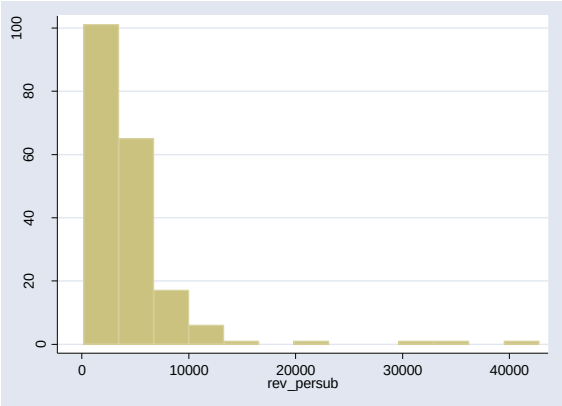

Figure 54 - Histogram of rev_persub (2005-2008)

H_0: botnet_activty and rev_persub are independent (ρ=0)
H_α: botnet_activty and rev_persub are positively correlated (ρ>0)

```
. spearman src persub rev persub
Number of obs =     194
Spearman's rho =       0.0788
Test of Ho: src persub and rev persub are independent
    Prob > |t| =       0.2746

. spearman spam_persub rev_persub
Number of obs =     194
Spearman's rho =      -0.0211
Test of Ho: spam persub and rev persub are independent
    Prob > |t| =       0.7703
```

Similar to the 2007 data, the test results are statistically insignificant at the 0.05 level. The null hypothesis cannot be rejected, (i.e., ISP's service pricing and botnet activity levels are unrelated). Figure 53 presents the scatter plots.

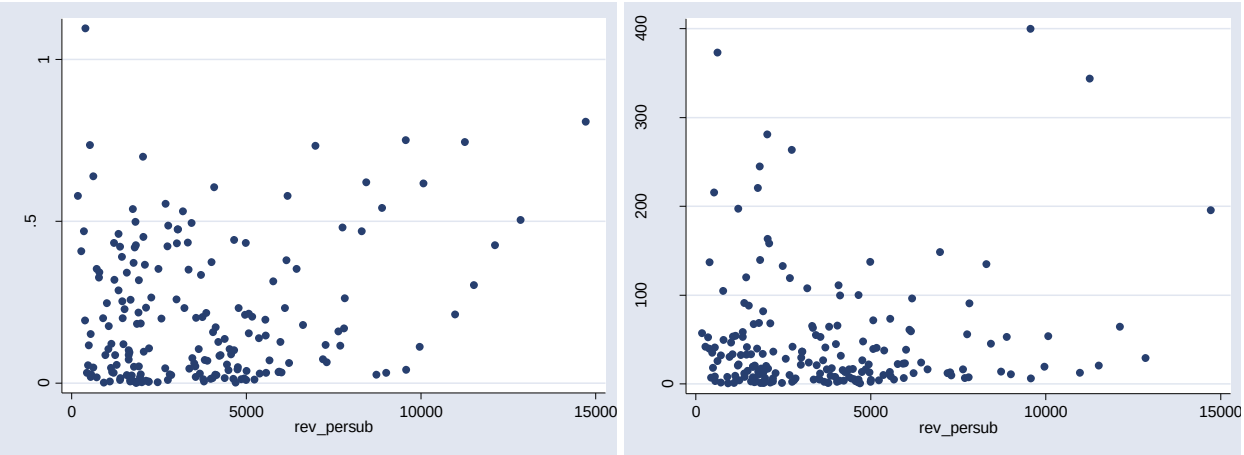

Figure 55 - Scatter plots: src_persub / spam_persub vs. rev_persub (2005-2008) – outliers removed

Finding

Based on all the above findings, the fourth hypothesis is rejected. That is, no relation is found to exist between an ISP's average revenue per user and its botnet activity levels. Caution must be taken into account when interpreting these results however. Not all operators report their revenue (the number of observations is approximately one third of the dataset), and for those that do, not all their revenue is attributable to retail broadband Internet services. These worries can be seen by value of *rev_persub* over $5000, which is quite doubtful (the average monthly price of broadband in the OECD would most probably be in the $20-$30 range).

5.2.5 HYPOTHESIS 5: CABLE PROVIDERS VS. DSL PROVIDERS

Cable providers have higher security performance than DSL providers.

Test strategy

We will divide the sample in two subsets, based on the value of *srv_type* (whether an ISP provides cable access or not). We will then perform a comparison of means on the variable representing botnet activity levels (*src_persub* or *spam_persub*) between these sub-groups. The actual statistical test can be either a ***t-test*** or a ***Wilcoxon rank sum test***, based on the assumptions of normality.[1]

2007 DATA

The *srv_type* variable has the values shown in Figure 56.

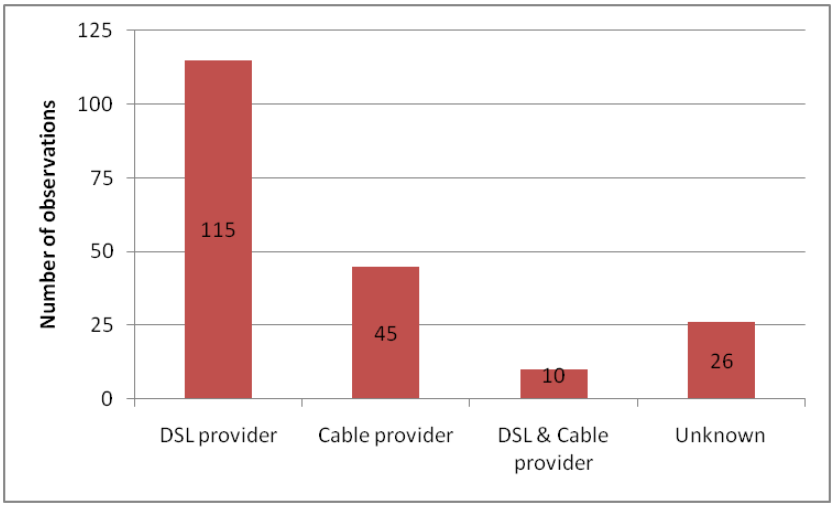

Figure 56 - Type of Internet access provided by ISPs (2007)

Test results

A new grouping variable, *srv_cable,* is created based on *srv_type*. Our two samples will be of size 115 and 55[2]. Although the variable to be tested (*src_persub* or *spam_persub*) is not from a normal distribution, due to the sample sizes being larger then 30, we can still use the parametric *t-test*[3]. The null and alternate hypotheses are as follows:

H_0: *The mean botnet_activty is equal in both mentioned groups of ISPs* ($\mu_{dsl} = \mu_{cable}$)

H_α: *The mean botnet_activty is smaller in ISPs providing cable access* ($\mu_{dsl} > \mu_{cable}$)

[1] Please note that *srv_type* has three levels: DSL, cable, and both. A variation of our strategy would be to create three subsets and compare the means between those using one-way ANOVA or Kruskal-Wallis test. That would separate ISPs that are providing both DSL & cable from those that are only providing cable; however, as our hypothesis is currently worded, we do not need to separate these last two groups from each other.

[2] Please note that we have included the hybrid providers to the cable providers in grouping, as this better matches the wording of the hypothesis (after all, the hybrids are providing cable); and as our theory of cable providers having monitoring equipment useful for botnet mitigation would hold for the hybrid providers too. Nonetheless, for safety, the statistical tests were rerun excluding the hybrid providers, and the same results were found.

[3] This is according to the central limit theorem

The box below contains the results of the t-test. (Before running the t-test, the variances of both samples were checked and found to be approximately equal). The result of the test is significant at the 0.05 level, and shows that the mean of *src_persub* (≈botnet activity level) is nearly 10% lower in ISPs that provide cable access.

```
. sdtest src_persub , by (srv_cable)
Variance ratio test
------------------------------------------------------------------------------
                            Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   1.4071                F =   1.4071               F =   1.4071
    P < F =   0.9189        2*(P > F) =   0.1623           P > F =   0.0811

. ttest src_persub , by (srv_cable)
Two-sample t test with equal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     115    .2611713    .0226541    .2429379    .2162938    .3060488
       1 |      55    .1611873    .0276155    .2048017    .1058216    .2165529
---------+--------------------------------------------------------------------
combined |     170    .2288235    .0180546    .2354028     .193182    .2644651
---------+--------------------------------------------------------------------
    diff |            .099984     .037931                .0251012    .1748669
------------------------------------------------------------------------------
Degrees of freedom: 168
                    Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   2.6359                 t =   2.6359                  t =   2.6359
   P < t =   0.9954             P > |t| =   0.0092              P > t =   0.0046
```

If we run these tests for our other variable, *spam_persub,* we paradoxically arrive at the opposite conclusion, and the outcome of the means comparison on the spam_persub variable is insignificant:

```
. sdtest spam_persub , by (srv_cable)
Variance ratio test
------------------------------------------------------------------------------
                            Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   0.8260                F =   0.8260               F =   0.8260
    P < F =   0.1970        2*(P < F) =   0.3939           P > F =   0.8030

. ttest spam_persub , by (srv_cable)
Two-sample t test with equal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     115    65.37478    6.055789     64.9411    53.37831    77.37125
       1 |      55    69.96309    9.635149    71.45618    50.64578     89.2804
---------+--------------------------------------------------------------------
combined |     170    66.85924    5.134067    66.94005    56.72407     76.9944
---------+--------------------------------------------------------------------
    diff |           -4.588308    11.00131               -26.30692    17.13031
------------------------------------------------------------------------------
Degrees of freedom: 168
                    Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =  -0.4171                 t =  -0.4171                  t =  -0.4171
   P < t =   0.3386             P > |t| =   0.6772              P > t =   0.6614
```

This difference might be explained by the fact that cable providers usually provide higher bandwidth to their subscribers. This means that despite having a lower percentage of bots on the networks, each infected machine sends out more spam, diminishing the effect, and leaving the average spam per subscriber unchanged. The spread of the tested variables can be visually compared in Figure 57.

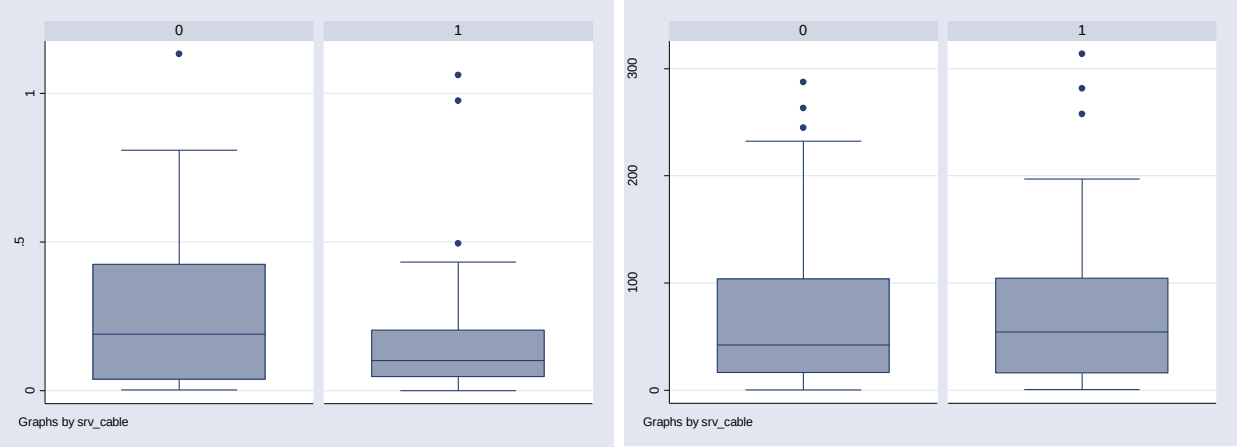

Figure 57 - Box plots of src_persub (left) and spam_persub (right), grouped by srv_cable (2007)

POOLED DATA

The *srv_type* variable has the values shown in Figure 58.

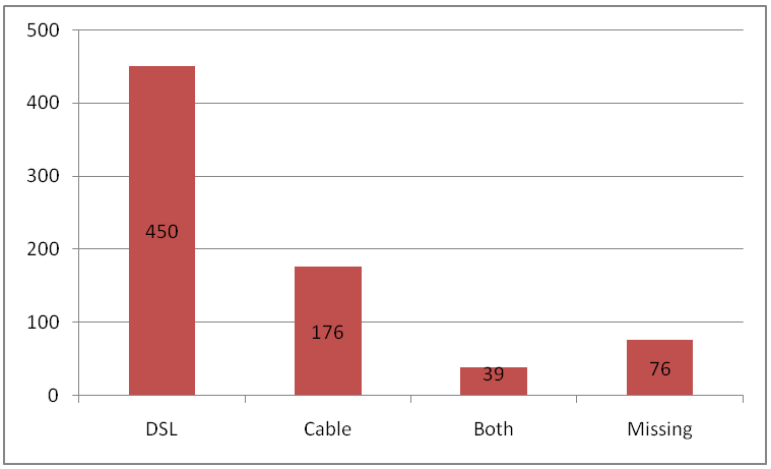

Figure 58 - Type of Internet access provided by ISPs (2005-2008)

Test results

Using the grouping variable *srv_cable,* we obtain two samples of size 450 and 215. Due to the sample sizes being ≥30, we can use the parametric *t-test*. The null and alternate hypotheses are as follows:

H_0: *The mean botnet_activty is equal in both mentioned groups of ISPs* ($\mu_{dsl} = \mu_{cable}$)
H_α: *The mean botnet_activty is smaller in ISPs providing cable access* ($\mu_{dsl} > \mu_{cable}$)

The box below contains the results of the t-test, and the prerequisite F-test (note that the equality of variances is rejected). The result of the t-test is significant at the 0.05 level, with the mean difference (of botnet activity levels) being a modest 8% lower in ISPs that provide cable access.

```
. sdtest src_persub , by (srv_cable)
Variance ratio test
------------------------------------------------------------------------------
                            Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)         Ha: sd(0) > sd(1)
        F =   1.5707                F =   1.5707              F =   1.5707
    P < F =   0.9999        2*(P > F) =   0.0002          P > F =   0.0001


. ttest src_persub , by (srv_cable) unequal
Two-sample t test with unequal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     450    .2095751    .0098228    .2083737    .1902707    .2288795
       1 |     215    .1329358    .0113389    .1662611    .1105856    .1552861
---------+--------------------------------------------------------------------
combined |     665     .184797    .0077125    .1988881    .1696531    .1999409
---------+--------------------------------------------------------------------
    diff |            .0766393     .015002                .047167    .1061116
------------------------------------------------------------------------------
Satterthwaite's degrees of freedom:  516.957
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   5.1086                 t =   5.1086                  t =   5.1086
   P < t =   1.0000           P > |t| =   0.0000            P > t =   0.0000
```

For the variable, *spam_persub,* we arrive at the similar result as the 2007 data - that is, we do not find a significant result at the 0.05 level. The explanation for this conflicting result would be similar to what was already mentioned. The box plots are presented in Figure 59.

```
. sdtest spam_persub , by (srv_cable)
Variance ratio test
------------------------------------------------------------------------------
                            Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)         Ha: sd(0) > sd(1)
        F =   1.1462                F =   1.1462              F =   1.1462
    P < F =   0.8723        2*(P > F) =   0.2554          P > F =   0.1277

. ttest spam_persub , by (srv_cable)
Two-sample t test with equal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     450    52.71723    3.443597    73.04973    45.94966     59.4848
       1 |     215    52.80613     4.65329    68.23063    43.63397    61.97828
---------+--------------------------------------------------------------------
combined |     665    52.74597    2.771715    71.47586    47.30359    58.18835
---------+--------------------------------------------------------------------
    diff |           -.0888989    5.930236               -11.7332    11.55541
------------------------------------------------------------------------------
Degrees of freedom: 663
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =  -0.0150                 t =  -0.0150                  t =  -0.0150
   P < t =   0.4940           P > |t| =   0.9880            P > t =   0.5060
```

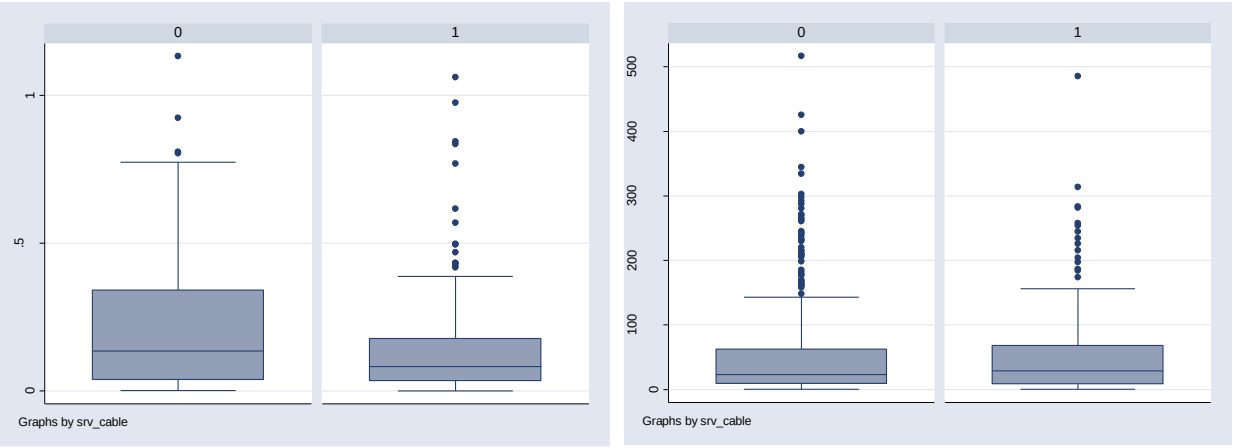

Figure 59 – Box plots of src_persub (right) and spam_persub (left), grouped by srv_cable (2005-2008)

Finding

The verdict for the fifth hypothesis depends on which of the two metrics for botnet activity we opt to use. Using *src_persub* as the metric, this hypothesis will be accepted; using *spam_persub,* it shall be rejected.

Based on theoretical discussions on the limitations of each dependent variable in Chapter 3, in this test *src_persub* would be a more robust indicator of botnet activity, as it measures whether a host is infected or not, irrelevant of the factors that influence the spamming capacity of the bot. This could well be the cause of the difference statistical results we are observing. Cable providers usually provide faster speeds, meaning that an infected bot can send out much more spam, so despite an 8% lower presence of bots on the networks of cable providers, the amount of outgoing spam (in relative terms) stay the same.

Unfortunately our data does not allow us to empirically test whether cable providers indeed provide faster speeds. Our nearest proxy is *int_bpp* (average Internet bandwidth per user, across the country the ISP operates in). This variable shows insignificant differences in means, when compared across the cable-providing and non-cable-providing ISPs, leaving us no choice but to sound a word of caution on our acceptance of this hypothesis.

5.2.6 HYPOTHESIS 6: EFFECTS OF REGULATION

ISPs in countries that have endorsed international agreements against cyber-crime (e.g., the LAP or the cyber-crime convention) have higher security performance.

Test strategy

We will run k-independent sample comparison of means, on the variables indicating botnet activity levels (*src_persub, spam_persub*). These samples will be constructed by grouping our observations by the regulatory framework that the ISP operates in - the variables *lap_mem* (2 levels) and *cyberconv_mem* (3 or 4 levels).

The statistical tests that can be used for this purpose differ based on the number of levels (subgroups), and whether a parametric or non-parametric version is required. We already know that our test variables are not normal. However, since n ≥ 30, for the two-level comparison we can still use the parametric ***t-test***. For the three-level comparison, no such equivalent rule exists, so the non-parametric ***Kruskal-Wallis one-way ANOVA test*** will be used.

2007 DATA

The *lap_mem* and *cyberconv_mem* variables have the frequencies presented below (n is ≥30 for all subgroups).

```
    lap mem |      Freq.     Percent
------------+----------------------------
          0 |         87       44.39
          1 |        109       55.61
------------+----------------------------
      Total |        196      100.00


cyberconv_m |
         em |      Freq.     Percent
------------+----------------------------
        non |         60       30.61
     signed |         86       43.88
   enforced |         50       25.51
------------+----------------------------
      Total |        196      100.00
```

Test results

First of all, let us run and present the results of the t-test, when grouped by *lap_mem*. The null and alternate hypotheses are as follows:

H_0: The mean botnet_activity is equal in both mentioned groups of ISPs ($\mu_{non\_lap} = \mu_{lap}$)
H_α: The mean botnet_activity is higher in ISPs operating in countries that are member of the LAP ($\mu_{non\_lap} > \mu_{lap}$)

The Stata outputs of the t-tests are presented in the boxes. As usual, the sample variances are checked before running t-tests.

```
. sdtest src_persub, by(lap_mem)
Variance ratio test
------------------------------------------------------------------------------
                             Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   1.2792                F =   1.2792               F =   1.2792
    P < F =   0.8875        2*(P > F) =   0.2250           P > F =   0.1125


. ttest src_persub, by(lap_mem)
Two-sample t test with equal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |      87    .3036989    .0260195    .2426939    .2519738    .3554239
       1 |     109    .1712615    .0205531    .2145806    .1305217    .2120013
---------+--------------------------------------------------------------------
combined |     196    .2300474    .0168774    .2362841    .1967617    .2633332
---------+--------------------------------------------------------------------
    diff |            .1324374    .0327027                .0679389    .1969359
------------------------------------------------------------------------------
Degrees of freedom: 194
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   4.0497                 t =   4.0497                  t =   4.0497
   P < t =   1.0000           P > |t| =   0.0001           P > t =   0.0000
```

```
. sdtest spam_persub, by(lap_mem)
Variance ratio test
------------------------------------------------------------------------------
                             Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   0.4750                F =   0.4750               F =   0.4750
    P < F =   0.0002        2*(P < F) =   0.0004           P > F =   0.9998


. ttest spam_persub, by(lap_mem) unequal
Two-sample t test with unequal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |      87    91.26437    8.414475    78.48499    74.53695    107.9918
       1 |     109     58.0033     10.9076    113.8787    36.38255    79.62405
---------+--------------------------------------------------------------------
combined |     196    72.76714    7.204602    100.8644     58.5582    86.97609
---------+--------------------------------------------------------------------
    diff |            33.26106    13.77603                6.087636    60.43449
------------------------------------------------------------------------------
Satterthwaite's degrees of freedom:    190.2
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   2.4144                 t =   2.4144                  t =   2.4144
   P < t =   0.9916           P > |t| =   0.0167           P > t =   0.0084
```

As can be seen, the null hypothesis can be rejected at the 0.05 significance level (using either measure of botnet activity). ISPs operating in countries that are a member of the London Action Plan have on average 13% less bot infections, and hence a higher security performance. This difference can be clearly seen in the box plots of Figure 60.

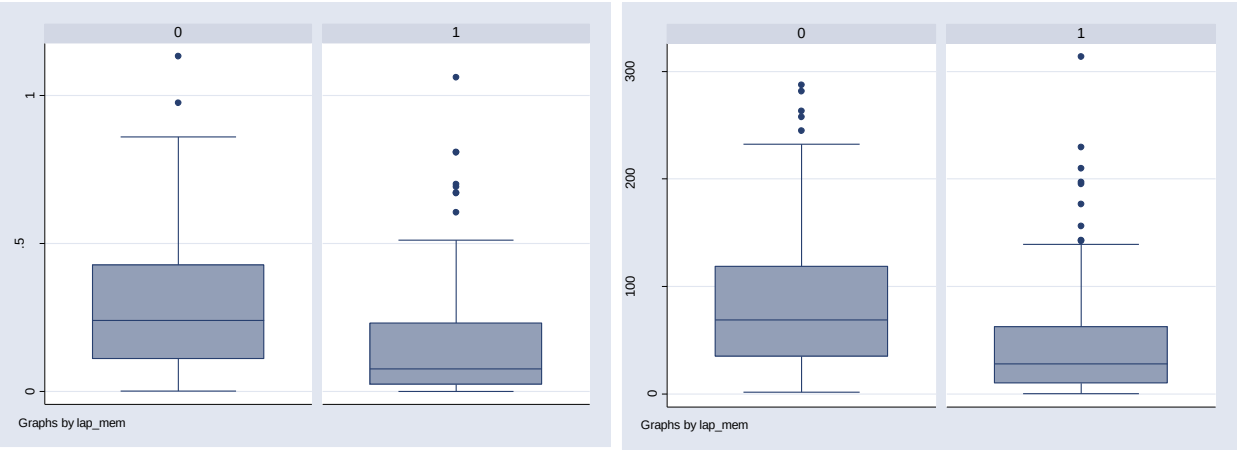

Figure 60 - Box plot of src_persub (left) and spam_persub (right), grouped by lap_mem (2007)

The next comparison of means we wish to do is when our observations are grouped by *cyberconv_mem*. As explained, since this variable has three levels (*non-member, signed, enforced*), and since our test variables (*src_persub* and *spam_persub*) are not normally distributed, we deploy the non-parametric Kruskal-Wallis one-way ANOVA test. The null and alternate hypotheses are as follows, and the results are presented after:

H_0: The mean botnet_activity is equal in all the mentioned groups of ISPs ($\mu_{non} = \mu_{signed} = \mu_{enforced}$)
H_α: The mean botnet_activity is not equal in all the mentioned groups of ISPs ($\mu_{non} \neq \mu_{signed}$ or $\mu_{signed} \neq \mu_{enforced}$)

```
. kwallis src persub, by(cyberconv mem)
Test: Equality of populations (Kruskal-Wallis test)
  +---------------------------+
  | cyberc~m | Obs | Rank Sum |
  |----------+-----+----------|
  |      non |  60 |  7016.00 |
  |   signed |  86 |  8331.50 |
  | enforced |  50 |  3958.50 |
  +---------------------------+
chi-squared =    12.213 with 2 d.f.
probability =     0.0022
chi-squared with ties =    12.213 with 2 d.f.
probability =     0.0022

. kwallis spam_persub, by(cyberconv_mem)
Test: Equality of populations (Kruskal-Wallis test)
  +---------------------------+
  | cyberc~m | Obs | Rank Sum |
  |----------+-----+----------|
  |      non |  60 |  6861.00 |
  |   signed |  86 |  7638.00 |
  | enforced |  50 |  4807.00 |
  +---------------------------+
chi-squared =     7.279 with 2 d.f.
probability =     0.0263
chi-squared with ties =     7.279 with 2 d.f.
probability =     0.0263
```

It can be seen that at the 0.05 significance level, the null hypothesis is rejected. The Kruskal-Wallis test does not however tell us which of the samples has a higher means. For this, we would need to do pair-wise comparisons of means. The results of the Wilcoxon ranksum test[1], presented below, show that at the 0.05 level, ISPs that

[1] We have used this test instead of t-test to avoid the extra prerequisite step of comparing variances.

operating in countries that have signed (or enforced) the Cybercrime convention, perform better in terms of security than the ISPs operating in non-member countries. (Please note that the resulting p-values need to be adjusted with Bonferroni's or similar method, to avoid type-I errors that occur in multiple comparison of means. The stated conclusion still holds after this adjustment[1])

```
. ranksum src_persub if cyberconv_mem != 3, by (cyberconv_mem)
Two-sample Wilcoxon rank-sum (Mann-Whitney) test
cyberconv_~m |      obs    rank sum    expected
-------------+---------------------------------
         non |       60      4946.5        4410
      signed |       86      5784.5        6321
-------------+---------------------------------
    combined |      146       10731       10731
Ho: src_pe~b(cyberc~m==non) = src_pe~b(cyberc~m==signed)
             z =   2.134
    Prob > |z| =   0.0328


. ranksum src_persub if cyberconv_mem != 0, by (cyberconv_mem)
Two-sample Wilcoxon rank-sum (Mann-Whitney) test
cyberconv_~m |      obs    rank sum    expected
-------------+---------------------------------
      signed |       86        6288        5891
    enforced |       50        3028        3425
-------------+---------------------------------
    combined |      136        9316        9316
Ho: src_pe~b(cyberc~m==signed) = src_pe~b(cyberc~m==enforced)
             z =   1.792
    Prob > |z| =   0.0732


. ranksum spam_persub if cyberconv_mem != 3, by (cyberconv_mem)
Two-sample Wilcoxon rank-sum (Mann-Whitney) test
cyberconv_~m |      obs    rank sum    expected
-------------+---------------------------------
         non |       60        5061        4410
      signed |       86        5670        6321
-------------+---------------------------------
    combined |      146       10731       10731
Ho: spam_p~b(cyberc~m==non) = spam_p~b(cyberc~m==signed)
             z =   2.589
    Prob > |z| =   0.0096


. ranksum spam_persub if cyberconv_mem != 0, by (cyberconv_mem)
Two-sample Wilcoxon rank-sum (Mann-Whitney) test
cyberconv_~m |      obs    rank sum    expected
-------------+---------------------------------
      signed |       86        5709        5891
    enforced |       50        3607        3425
-------------+---------------------------------
    combined |      136        9316        9316
Ho: spam_p~b(cyberc~m==signed) = spam_p~b(cyberc~m==enforced)
             z =  -0.821
    Prob > |z| =   0.4114
```

The fact that the spread in botnet activity is larger in ISPs operating in non-member countries can also be seen in the box plots of Figure 61.

[1] The formula for Bonferonni's method is: *p' = min(1,np)*

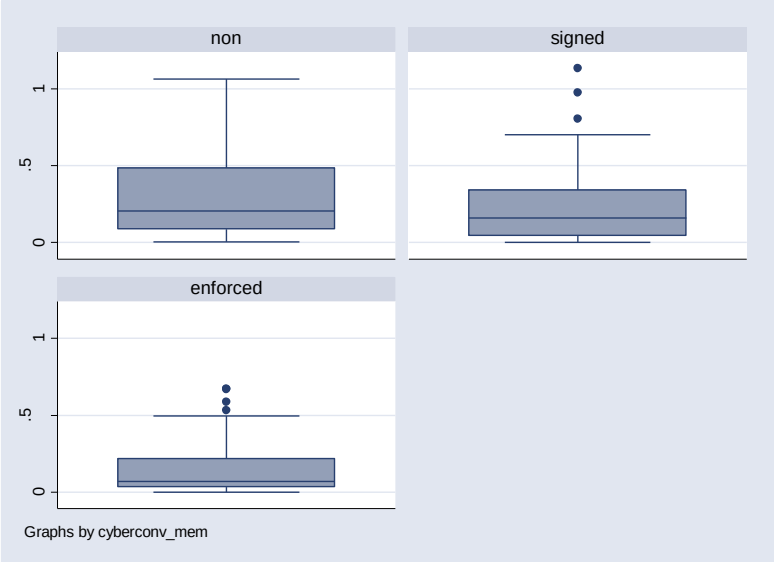

Figure 61 - Box plot of src_persub, grouped by cyberconv_mem (2007)

POOLED DATA

The procedure for the pooled dataset is similar to what was explained for the 2007 dataset. The *lap_mem* and *cyberconv_mem* variables have the frequencies presented below.

```
    lap_mem |      Freq.     Percent
------------+----------------------------
          0 |        324       43.72
          1 |        417       56.28
------------+----------------------------
      Total |        741      100.00


cyberconv_m |
         em |      Freq.     Percent
------------+----------------------------
 non-member |        215       29.01
     signed |        366       49.39
   ratified |         20        2.70
   enforced |        140       18.89
------------+----------------------------
      Total |        741      100.00
```

Test results

First of all, let us run and present the results of the t-test, when grouped by *lap_mem*. The null and alternate hypotheses are as follows:

H_0: The mean botnet_activity is equal in both mentioned groups of ISPs ($\mu_{non\_lap} = \mu_{lap}$)
H_α: The mean botnet_activity is higher in ISPs operating in countries that are member of the LAP ($\mu_{non\_lap} > \mu_{lap}$)

```
. sdtest src_persub, by(lap_mem)
Variance ratio test
------------------------------------------------------------------------------
                              Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   1.3584                F =   1.3584               F =   1.3584
    P < F =   0.9983        2*(P > F) =   0.0033           P > F =   0.0017

. ttest src_persub, by(lap_mem) unequal
Two-sample t test with unequal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     324    .2587201    .0121807    .2192527    .2347565    .2826836
       1 |     417    .1390209     .009212    .1881154    .1209129    .1571288
---------+--------------------------------------------------------------------
combined |     741     .191359    .0077413    .2107284    .1761614    .2065565
---------+--------------------------------------------------------------------
    diff |            .1196992    .0152719                .0897098    .1496886
------------------------------------------------------------------------------
Satterthwaite's degrees of freedom:  636.484
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   7.8379                 t =   7.8379                  t =   7.8379
   P < t =   1.0000           P > |t| =   0.0000           P > t =   0.0000
```

```
. sdtest spam persub, by(lap mem)
Variance ratio test
------------------------------------------------------------------------------
                              Ho: sd(0) = sd(1)
   Ha: sd(0) < sd(1)          Ha: sd(0) != sd(1)          Ha: sd(0) > sd(1)
        F =   1.9500                F =   1.9500               F =   1.9500
    P < F =   1.0000        2*(P > F) =   0.0000           P > F =   0.0000


. ttest spam_persub, by(lap_mem) unequal
Two-sample t test with unequal variances
------------------------------------------------------------------------------
   Group |     Obs        Mean    Std. Err.   Std. Dev.   [95% Conf. Interval]
---------+--------------------------------------------------------------------
       0 |     324    74.58143    5.040368    90.72662    64.66533    84.49752
       1 |     417     40.8034    3.181665    64.97145    34.54925    47.05754
---------+--------------------------------------------------------------------
combined |     741    55.57274     2.90354    79.03818    49.87258     61.2729
---------+--------------------------------------------------------------------
    diff |            33.77803    5.960562                22.07034    45.48571
------------------------------------------------------------------------------
Satterthwaite's degrees of freedom:  562.362
                     Ho: mean(0) - mean(1) = diff = 0
     Ha: diff < 0                 Ha: diff != 0                 Ha: diff > 0
       t =   5.6669                 t =   5.6669                  t =   5.6669
   P < t =   1.0000           P > |t| =   0.0000           P > t =   0.0000
```

As can be seen, the null hypothesis (for both measures of botnet activity) is rejected at the 0.05 significance level. (ISPs operating in countries that are member of the London Action Plan have on average 12% less bots.) The box-plots in Figure 62 illustrate this point.

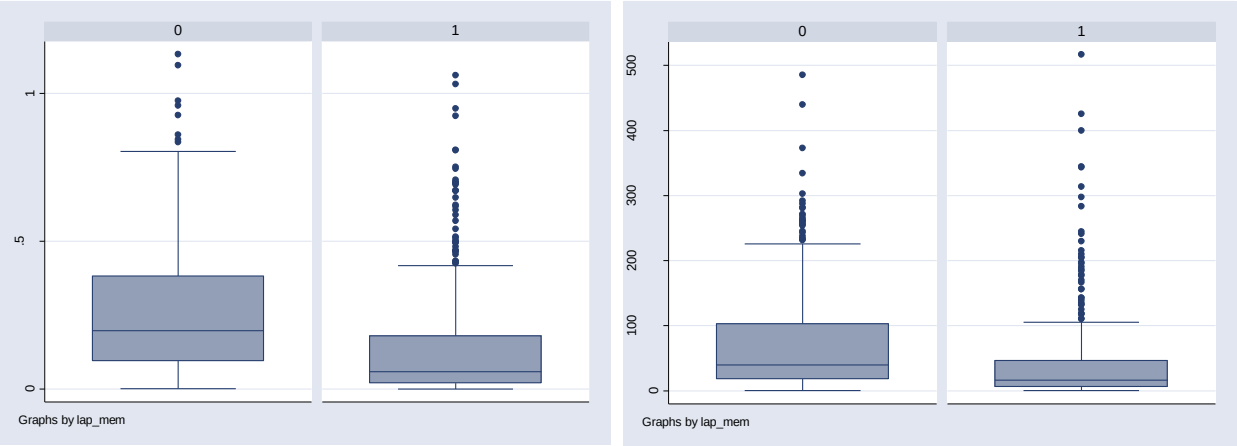

Figure 62 - Box plot of src_persub (left) and spam_persub (right), grouped by lap_mem (2005-2008)

The next comparison of means is performed by grouping by *cyberconv_mem*. In the pooled dataset, this variable has four levels (*non-member, signed, ratified, enforced*). We will again deploy the Kruskal-Wallis one-way ANOVA test. The null and alternate hypotheses are as follows, and the results are presented in the box beneath it:

H_0: The mean botnet_activity is equal in all the mentioned groups of ISPs ($\mu_{non} = \mu_{signed} = \mu_{ratified} = \mu_{enforced}$)
H_α: The mean botnet_activity is not equal in all the mentioned groups of ISPs

```
. kwallis src_persub, by(cyberconv_mem)
Test: Equality of populations (Kruskal-Wallis test)
  +---------------------------+
  | cyberc~m | Obs |  Rank Sum |
  |----------+-----+-----------|
  | non-memb | 215 |  99067.50 |
  |   signed | 366 | 126367.00 |
  | ratified |  20 |   5136.50 |
  | enforced | 140 |  44340.00 |
  +---------------------------+
chi-squared =    57.807 with 3 d.f.
probability =     0.0001
probability =     0.0001

. kwallis spam_persub, by(cyberconv_mem)
Test: Equality of populations (Kruskal-Wallis test)
  +---------------------------+
  | cyberc~m | Obs |  Rank Sum |
  |----------+-----+-----------|
  | non-memb | 215 |  90858.00 |
  |   signed | 366 | 119738.00 |
  | ratified |  20 |   4443.00 |
  | enforced | 140 |  59872.00 |
  +---------------------------+
chi-squared =    47.329 with 3 d.f.
probability =     0.0001
probability =     0.0001
```

It can be seen that at the 0.05 significance level, the null hypothesis is rejected. As mentioned, the Kruskal-Wallis test does not tell us which of the samples has a higher means. We would need to do multiple pair-wise comparisons of means. However, the number of pair-wise permutations is considerably high – (4!/2!2! = 6), multiplied by the two test variables. We will instead opt to perform the comparisons visually with the aid of box-plots, presented in Figure 63 and Figure 64. The biggest difference is in the non-member group versus the others.

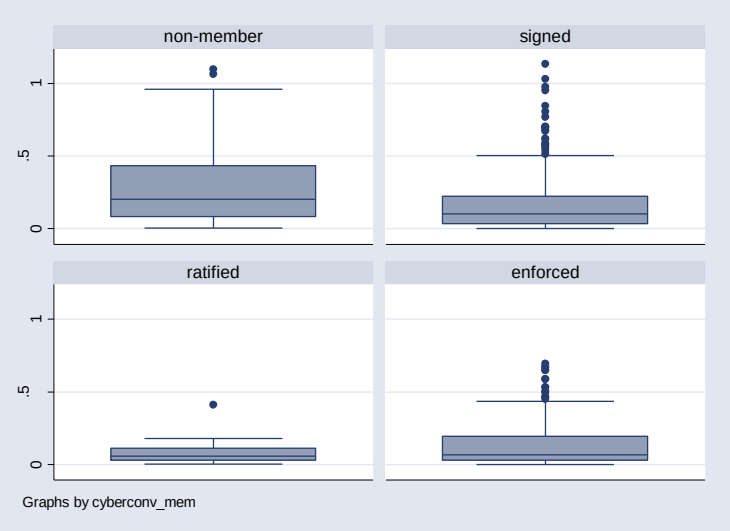

Figure 63 - Box plot of src_persub, grouped by cyberconv_mem (2005-2008)

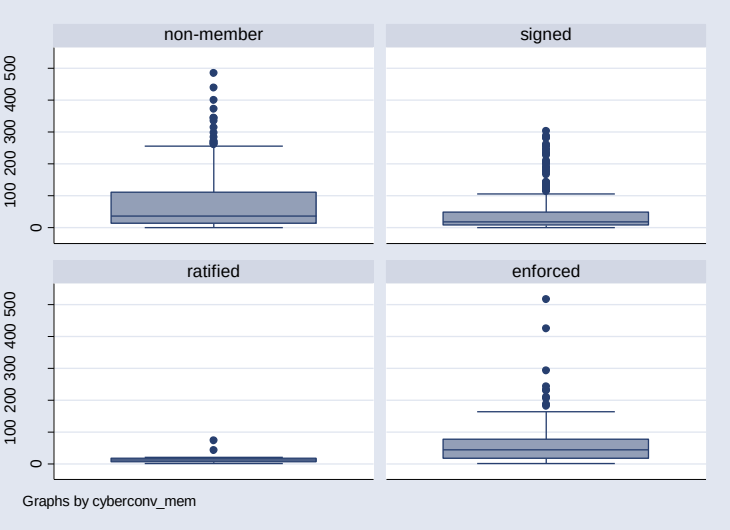

Figure 64 - Box plot of spam_persub, grouped by cyberconv_mem (2005-2008)

Finding

Based on the empirical evidence and statistical tests performed, we can accept the sixth hypothesis. ISPs that operate in countries that are member of the London Action Plan (which account for around half of our observations), have on average 10% lower number of bots. The difference is moderate, yet statistically significant at the 0.05 level. A same statement can be made about countries that have signed the Cybercrime convention.[1]

The acceptance of this hypothesis can be good news for policy makers, as it appears to show that regulation can help tackle the botnet problem, albeit moderately. However, we must take care not to draw premature conclusions, as we can only be certain that a relation exists, but cannot say anything about the causality. It could be possible that countries that adopt these treaties have a more educated population, and in fact the higher education of end-users is causing lower bot infections; alternatively, it could be that income in these countries is higher, and buying AV software is affordable for people; etc. We will return to these points in the conclusion.

[1] Note that signing the treaty matters - no significant differences between countries enforcing the treaty and those only signing

5.2.7 HYPOTHESIS 7: EFFECTS OF PIRACY

ISPs in countries with higher piracy rates have lower security performance.

Test strategy

Statistical tests to measure the degree of association between the following variables needs to be used:

- *src_persub* or *spam_persub* (botnet activity level / security performance)
- *piracy_rate*

For this purpose, the non-parametric ***Spearman rank correlation coefficient tests*** will be used (as we already know that the assumption of normality does not hold). A scatter plot will also be drawn.

2007 data

The summary statistics, histograms, and test of normality for the independent variable *piracy_rate* are as provided below. The distribution for this variable is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
 piracy rate |       196    40.68367    17.75739         20         84

                  Skewness/Kurtosis tests for Normality
                                                       ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+-------------------------------------------------------
 piracy_rate |      0.000         0.055          17.06       0.0002
```

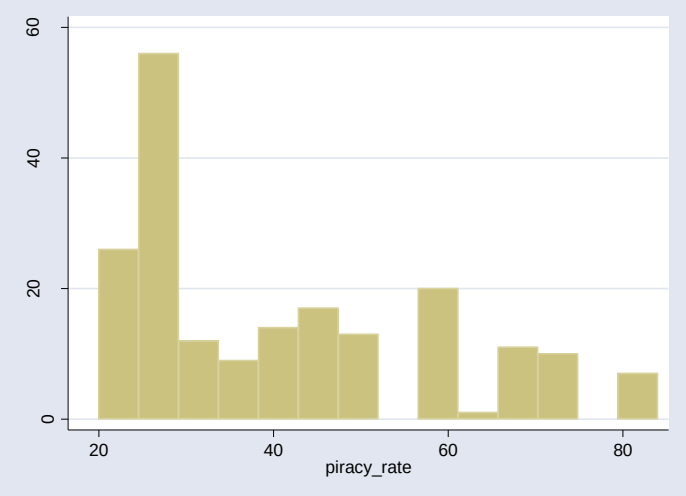

Figure 65 - Histogram of piracy_rate (2007)

Test results

Due to the lack of normality, the non-parametric tests to measure the degree of association are used. The null and alternate hypotheses for these tests, and the actual test results, are as follows:

H_0: botnet_activity and piracy_rate are independent (ρ=0)
H_α: botnet_activity and piracy_rate are positively correlated (ρ>0)

```
. spearman piracy rate src persub
Number of obs =     196
Spearman's rho =       0.3500
Test of Ho: piracy_rate and src_persub are independent
    Prob > |t| =       0.0000


. spearman piracy_rate spam_persub
Number of obs =     196
Spearman's rho =       0.4527
Test of Ho: piracy_rate and spam_persub are independent
    Prob > |t| =       0.0000
```

The results of both rank correlation tests show that at the 0.05 significance level, the null hypothesis can be rejected. In other words, we can assume that piracy rate is correlated with ISP botnet activity levels. The direction of the association is positive, meaning that ISPs performing in markets with higher piracy rates have higher botnet activity levels. Please note that the degree of association is moderate (ρ = 0.35).

The scatter plots in Figure 66 depict these relations graphically. At first glance, the scatter-plots appear to contradict the test results, as they show dots all over the place. The identified association can be understood in the graphs by the cluster in bottom-left corner - a whole group of ISPs with low botnet activity levels, operating in countries with low piracy rates. The spread on the right side of the graphs is larger. The association exists, but is not very strong. [1]

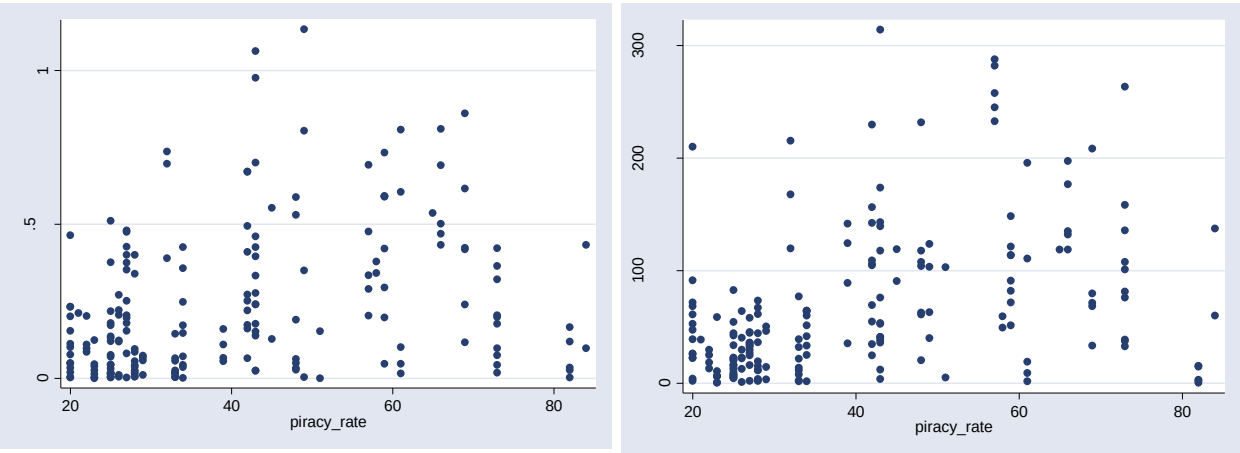

Figure 66 - Scatter plots: left: src_persub vs. piracy_rate, right: spam_persub vs. piracy_rate (2007, outliers removed)

Pooled data

The summary statistics, histograms, and test of normality for the independent variable *piracy_rate* are as follows. The distribution for this variable is not normal, so non-parametric tests need to be used.

[1] Please do not forget that we have used a 'rank correlation test' – so you shouldn't expect to see a linear relation in the graphs, but rather a more generic *monotone* function.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
 piracy_rate |       740    40.35135    17.31936         20         87


                  Skewness/Kurtosis tests for Normality
                                                       ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)     Prob>chi2
-------------+-------------------------------------------------------------
 piracy_rate |      0.000         0.012           57.46        0.0000
```

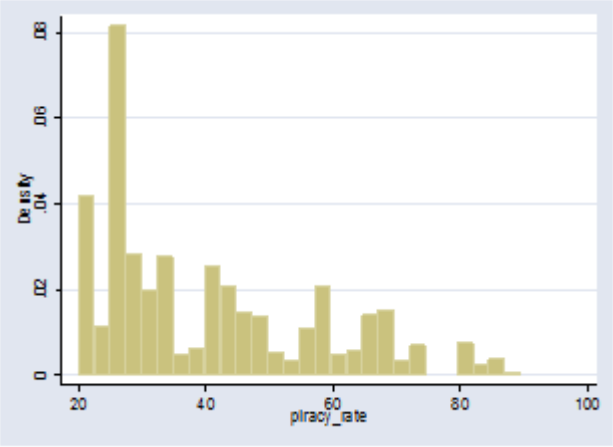

Figure 67 - Histogram of piracy_rate (2005-2008)

Test results

The null and alternate hypotheses for these tests, and the actual test results, are as follows:

H_0: botnet_activity and piracy_rate are independent (ρ=0)
H_α: botnet_activity and piracy_rate are positively correlated (ρ>0)

```
. spearman piracy_rate src_persub
Number of obs =     740
Spearman's rho =       0.3910
Test of Ho: piracy_rate and src_persub are independent
    Prob > |t| =       0.0000

. spearman piracy_rate spam_persub
Number of obs =     740
Spearman's rho =       0.3424
Test of Ho: piracy_rate and spam_persub are independent
    Prob > |t| =       0.0000
```

Similar to the 2007 data, the results of the rank correlation tests show that at the 0.05 significance level, the null hypothesis can be rejected. We can assume that piracy rate is correlated with ISP botnet activity levels. The direction of the association is positive, and moderate (ρ = 0.39). The scatter plots in Figure 66 illustrate these relations graphically. Again, although the scatter plots show the dots scattered all over, a bottom-left cluster of ISPs exists – which have low botnet activity levels and operate in countries with low piracy rates.

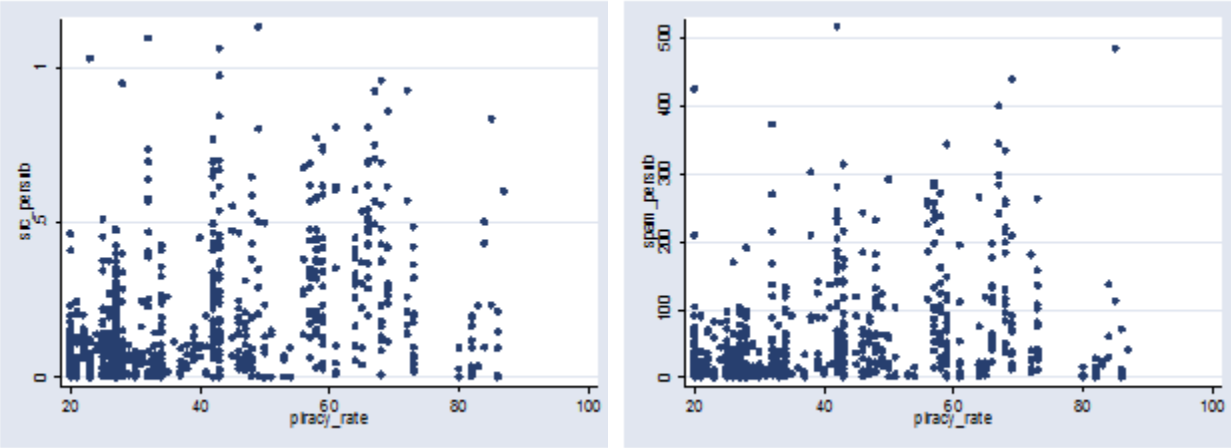

Figure 68 - Scatter plots: src_persub and spam_persub, vs. piracy_rate (2005-2008) – outliers removed

Finding

Based on the results of the statistical tests, we can accept our seventh hypothesis. That is, lower piracy rate is associated with lower levels of botnet activity (both in terms of number of bot infections and spam messages emitted), at a moderate degree ($\rho \approx 0.35$).

However, exactly as we explained for the previous hypothesis, association does not equal causation. There might well be latent variables involved here. For instance, countries with lower piracy rates are quite likely to have become member of the LAP – the two variables are significantly associated with each other:

```
. spearman lap_mem piracy_rate
 Number of obs =     740
Spearman's rho =      -0.3740
Test of Ho: lap_mem and piracy_rate are independent
    Prob > |t| =       0.0000
```

This obviously means that factors other than piracy rates may well be at play. When in section 5.3 we move towards a full regression model, the interplay between these variables will be examined.

5.2.8 HYPOTHESIS 8: EFFECTS OF BANDWIDTH

ISPs in countries with higher avg. bandwidth rates have lower security performance.

Test strategy

Statistical tests to measure the degree of association between the following variables needs to be used:

- *src_persub* or *spam_persub* (botnet activity level / security performance)
- *int_bpp* (average Internet bandwidth per user, in bits per second)

For this purpose, the non-parametric ***Spearman rank correlation coefficient test*** will be used (as we already know that the assumption of normality does not hold.) A scatter plot will also be drawn.

2007 DATA

The summary statistics, histograms, and test of normality for the independent variable *int_bpp* are as follows. The distribution for this variable is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
     int_bpp |       195    20414.71    20448.75        441      92832


                  Skewness/Kurtosis tests for Normality
                                                 ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)     Prob>chi2
-------------+-------------------------------------------------------
     int_bpp |      0.000         0.000          52.43        0.0000
```

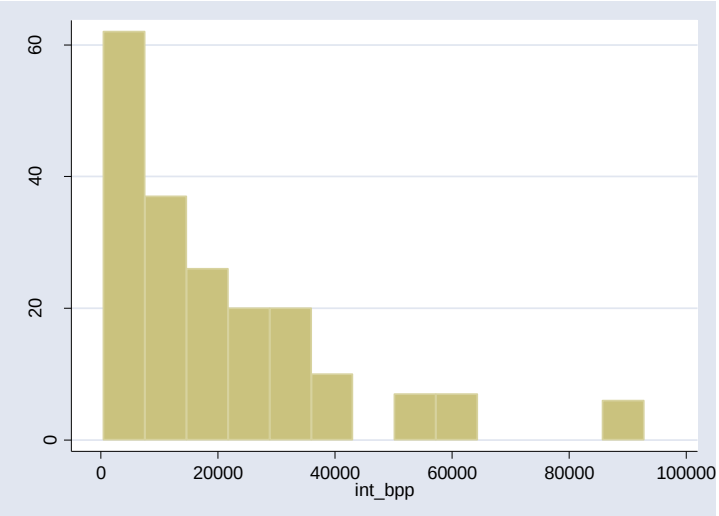

Figure 69 - Histogram of int_bpp (2007)

Test results

The null and alternate hypotheses for the non-parametric measure of association test, and the actual test results, are as follows:

H_0: botnet_activity and int_bpp are independent (ρ=0)
H_α: botnet_activity and int_bpp are positively associated (ρ >0)

```
. spearman int bpp src persub
Number of obs =     195
Spearman's rho =      -0.1977
Test of Ho: int_bpp and src_persub are independent
    Prob > |t| =       0.0056

. spearman int_bpp spam_persub
Number of obs =     195
Spearman's rho =      -0.2758
Test of Ho: int_bpp and spam_persub are independent
    Prob > |t| =       0.0001
```

The results show that at the 0.05 significance level, the null hypothesis can be rejected. We can assume that the average bandwidth available to each user is associated with ISP botnet activity levels. The association is rather weak (ρ = -0.20), and unexpectedly negative. The scatter plots in Figure 70 visualize these associations, with the now familiar bottom left cluster. We will give a possible explanation for the observation that botnet infections are higher in ISPs operating in countries with lower bandwidth, contrary to the literature, under 'Findings'.

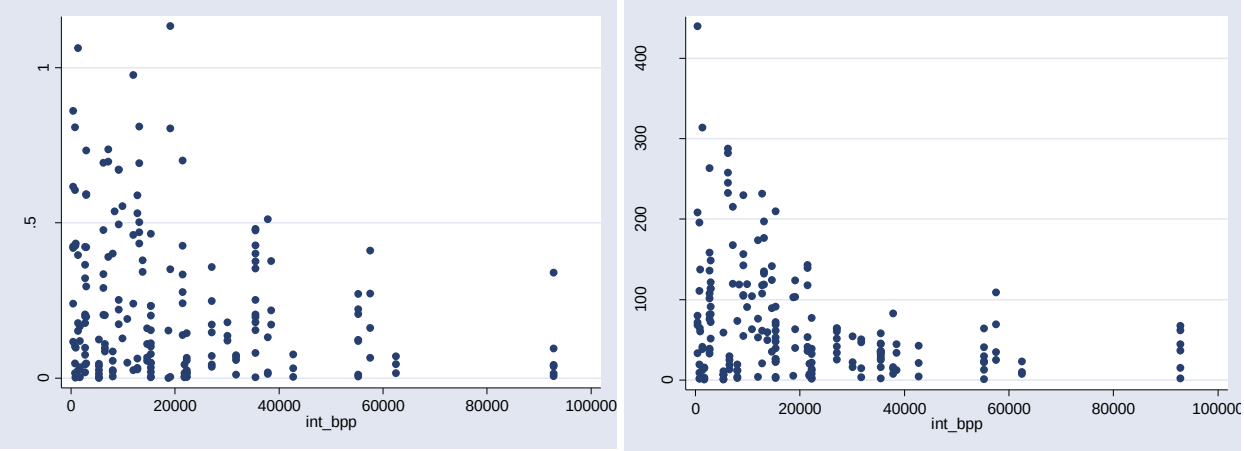

Figure 70 - Scatter plots: src_persub / spam_persub, versus int_bpp (2007)

POOLED DATA

The summary statistics, histograms, and test of normality for the independent variable *int_bpp* are as follows. The distribution for this variable is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
     int_bpp |       386    14345.77    16918.12   190.8559   92832.46
```

```
             Skewness/Kurtosis tests for Normality
                                                       ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+-------------------------------------------------------
     int_bpp |      0.000         0.000              .       0.0000
```

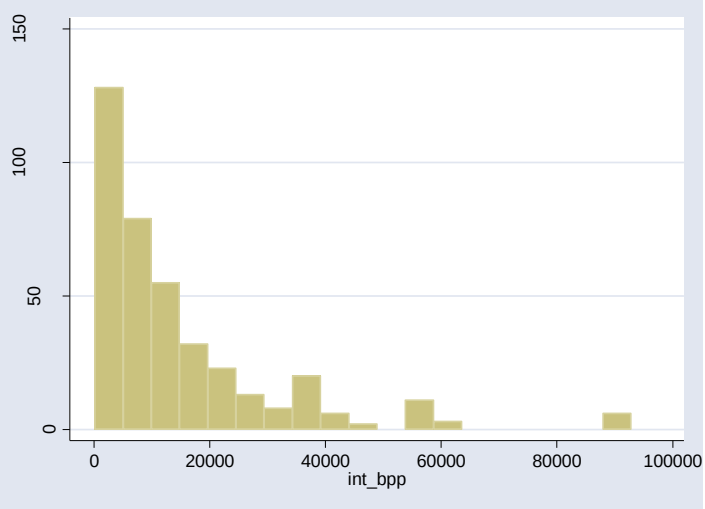

Figure 71 - Histogram of int_bpp (2005-2008)

Test results

The null and alternate hypotheses for the statistical tests, and the actual test results, are as follows:

H_0: botnet_activity and int_bpp are independent (ρ=0)
H_α: botnet_activity and int_bpp are positively associated (ρ >0)

```
. spearman int bpp src persub
Number of obs =     386
Spearman's rho =      -0.2324
Test of Ho: int_bpp and src_persub are independent
    Prob > |t| =         0.0000


. spearman int_bpp spam_persub
Number of obs =     386
Spearman's rho =      -0.0411
Test of Ho: int_bpp and spam_persub are independent
    Prob > |t| =         0.4209
```

The results show that at the 0.05 significance level, the null hypothesis can be rejected, when *src_persub* is used as the metric for botnet activity. In this case, we can assume that the average bandwidth available to each user is associated with ISP botnet activity levels. The association is rather weak (ρ=-0.23), and as we already saw for the 2007 data, has a negative direction. Figure 72 shows the relevant scatter plots.

Interestingly enough however, if *spam_persub* is used as the metric, the result is insignificant. This is similar to what we saw for the hypothesis involving cable providers: a decreased number of bot infections, which due to the higher bandwidth (capacity) available to these bots, leave the average spam sent per subscriber unaffected.

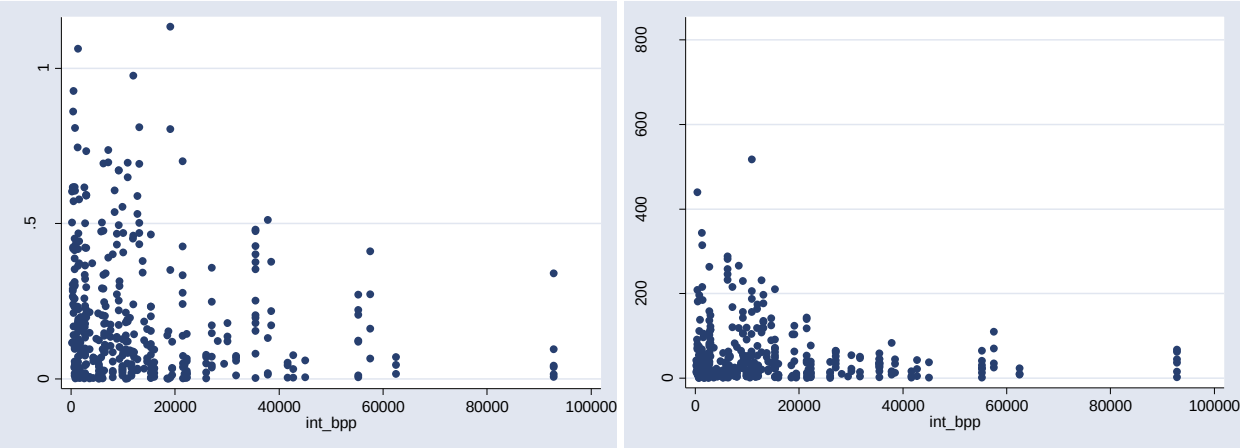

Figure 72 - Scatter plots: src_persub / spam_persub, versus int_bpp (2007)

Finding

Based on the results of the statistical tests, the eighth hypothesis is rejected. Although we do find an association between the average Internet bandwidth, of the country that an ISP operates in, and the botnet activity levels of the ISP. But the direction is opposite of what is hypothesized. 'Bandwidth' is recognized as one of the enablers of malware in the literature, so we would expect ISPs operating in countries with higher average Internet bandwidth to have more bot infections – which we are not. The answer may well lie in the fact that Internet access speed is highly correlated with broadband penetration rates, and in turn, with the economic infrastructure, GDP per capita, LAP membership, piracy rates, etc. For instance, the association between piracy rates and Internet speeds is quite strong (ρ = -0.68).

```
. spearman int_bpp piracy_rate
 Number of obs =     385
Spearman's rho =      -0.6802
Test of Ho: int_bpp and piracy_rate are independent
    Prob > |t| =       0.0000
```

To be able to truly assess these relations, we would need to control for the other variables, which we shall do in section 5.3 (full model). Having more fine-grained data on the bandwidth per ISP would also help, instead of looking at a country level average.

Despite this interplay limitation, these results can still be interpreted positively for policy makers: increased broadband penetration (and speeds) does not "automatically" translate into higher percentages of bot infections; and that there exists factors that diffuse the security side-effects of bandwidth

5.2.9 HYPOTHESIS 9: EFFECTS OF EDUCATION

ISPs in countries with a higher educational index have higher security performance.

Test strategy

Statistical tests to measure the degree of association between the following variables needs to be used:

- *src_persub* or *spam_persub* (botnet activity level / security performance)
- *educ_ix* (educational index)

For this purpose, the non-parametric ***Spearman rank correlation coefficient test*** will be used (knowing that the assumption of normality does not hold.) Scatter plots is also employed.

2007 DATA

The summary statistics, histograms, and test of normality for the independent variable *educ_ix* are as follows. The distribution for is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
     educ_ix |       196    .9442653    .0652522       .643       .993


                  Skewness/Kurtosis tests for Normality
                                                 ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)    Prob>chi2
-------------+-------------------------------------------------------
     educ_ix |      0.000         0.000              .       0.0000
```

Test results

The null and alternate hypotheses for Spearman's rank correlation coefficient test, and the actual test results, are as follows:

H_0: botnet_activity and educ_ix are independent (ρ=0)
H_α: botnet_activity and educ_ix are negatively associated (ρ<0)

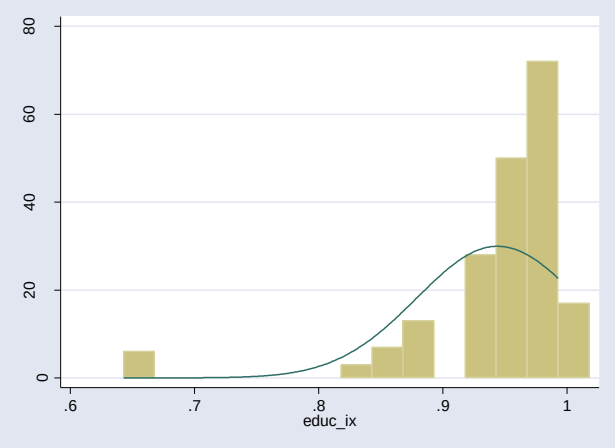

Figure 73 - Histogram of educ_ix (2007)

```
. spearman src_persub educ_ix
Number of obs =      196
Spearman's rho =      -0.3536
Test of Ho: src_persub and educ_ix are independent
    Prob > |t| =       0.0000

. spearman spam persub educ ix
Number of obs =      196
Spearman's rho =      -0.3011
Test of Ho: spam_persub and educ_ix are independent
    Prob > |t| =       0.0000
```

The results show that at the 0.05 significance level, the null hypothesis can be rejected. We can assume that the educational index of the country an ISP operates in, associates with the ISP's botnet activity levels. The correlation rank is negative (as expected), and moderately strong (ρ=-0.35).

The scatter plots in Figure 74 illustrate these relations graphically. The negative association can be understood by the large number of observations that have high *educ_ix* and low botnet activity, zipped together in the bottom right corner of the graph, versus the observations with lower *educ_ix* and a wider spread in botnet activity levels.

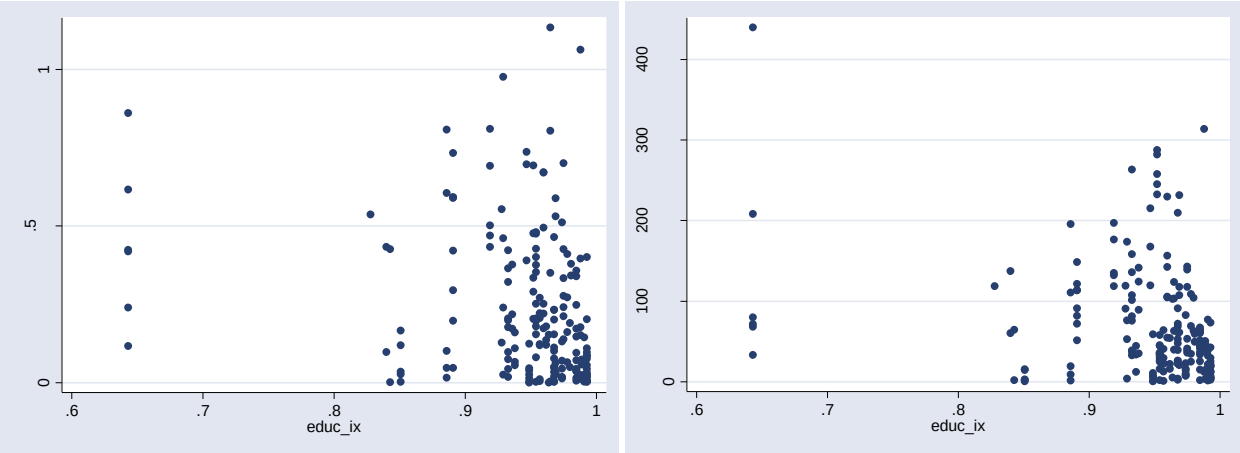

Figure 74- Scatter plots: src_persub / spam_persub, versus educ_ix (2007)

POOLED DATA

The summary statistics, histograms, and test of normality for the independent variable *educ_ix* are as follows. The distribution for this variable is not normal.

```
    Variable |       Obs        Mean    Std. Dev.       Min        Max
-------------+--------------------------------------------------------
     educ_ix |       547    .9439561    .0678809       .632       .993

                    Skewness/Kurtosis tests for Normality
                                                   ------- joint ------
    Variable |  Pr(Skewness)   Pr(Kurtosis)  adj chi2(2)     Prob>chi2
-------------+-------------------------------------------------------
     educ_ix |      0.000          0.000              .       0.0000
```

Test results

The null and alternate hypotheses, and the actual test results, are as follows:

H_0: botnet_activity and educ_ix are independent (ρ=0)

H_α: botnet_activity and educ_ix are negatively associated(ρ<0)

Similar to the 2007 data, the results are significant at the 0.05 level. The null hypothesis can be rejected. The rank correlation is negative (as expected), and moderate (ρ=-0.38). The scatter plots illustrate these associations.

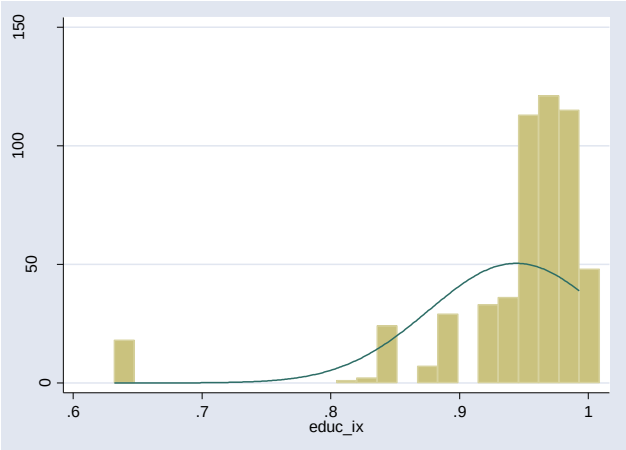

Figure 75 - Histogram of educ_ix (2005-2008)

```
. spearman src_persub educ_ix
Number of obs =      547
Spearman's rho =      -0.3805
Test of Ho: src_persub and educ_ix are independent
    Prob > |t| =       0.0000

. spearman spam_persub educ_ix
Number of obs =      547
Spearman's rho =      -0.2606
Test of Ho: spam_persub and educ_ix are independent
    Prob > |t| =       0.0000
```

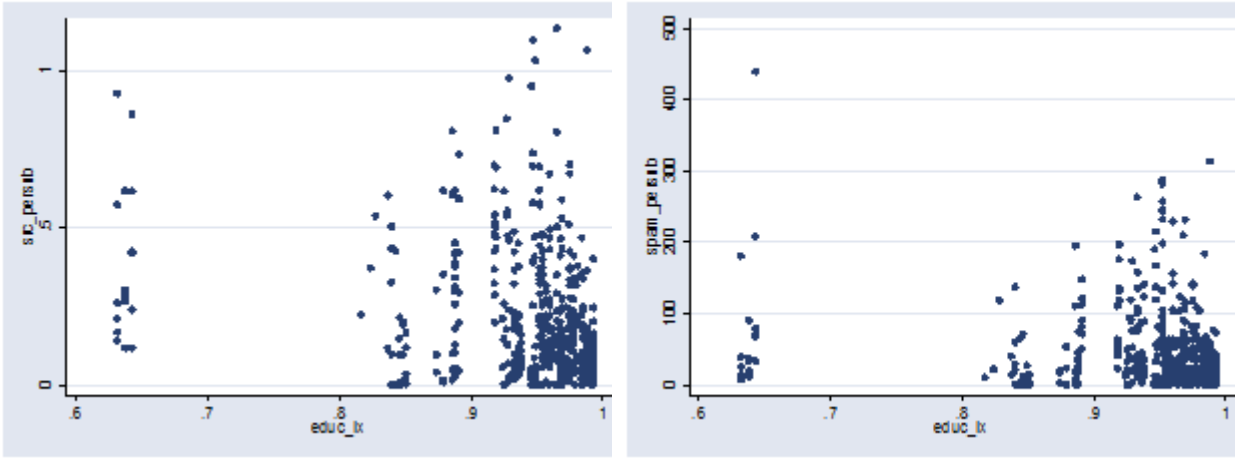

Figure 76 - Scatter plots: src_persub / spam_persub, versus educ_ix (2005-2008)

Finding

Based on the statistical test results, the ninth hypothesis is accepted. ISPs operating in countries with higher levels of education have lower levels of botnet activity. The now familiar warning must be raised that we are concluding mere associations between these variables, which when controlled for other variables might cease to exist.

5.2.10 SUMMARY

The following table lists the summary of the findings in this section. The table includes:

- the hypotheses,
- the statistical instrument used for testing the hypotheses,
- the four test results (two independent variables multiplied by two datasets);
- and finally, the verdict for each individual hypothesis.

Table 27 - Summary of statistical test results and findings for each hypothesis

#	Subject	Independent variables	Statistical instrument	N 07/pooled	Test results *07, srcs*	Test results *07, msgs*	Test results *pooled, srcs*	Test results *pooled,msgs*	Verdict
1	**ISPs are central**	-	Descriptive statistics	*N=RAW DATASET*	*ratio=82%*	*ratio=72%*	*ratio=81%*	*ratio=72%*	**Accepted**
2	**ISPs differ significantly**	-	Descriptive statistics	*N=196/741*	*CV = 1.03*	*CV=1.39*	*CV = 1.10*	*CV=1.42*	**Accepted**
3	**Effects of ISP size**	total_sub market_share	Spearman's rho	*N=196/741*	*sig_1 = 0.108* *sig_2 = 0.600*	*sig_1 = 0.000* *ρ = -0.261* *sig_2= .950*	*sig_1=0.000* *ρ = -0.170* *sig_2= .820*	*sig_1=0.000* *ρ = -0.182* *sig_2= .179*	**Rejected** (opposite holds with total_sub)
4	**Effects of ARPU**	rev_persub	Spearman's rho	*N=59/194*	*sig = 0.417*	*sig = 0.382*	*sig = 0.275*	*sig=0.770*	**Rejected**
5	**Cable providers vs. DSL providers**	srv_cable	t-test	*N=170/665*	*sig = 0.005* *diff = .0999*	*sig=0.661*	*sig = 0.000* *diff = .0766*	*sig = 0.506*	**Accepted** (with *srcs*, not *msgs*)
6	**Effects of regulation**	lap_mem cyberconv_mem	t-test, Kruskal-Wallis	*N=196/741*	*sig_1 = 0.000* *diff = 0.132* *sig_2 = 0.002*	*sig_1 = 0.008* *diff = 33.26* *sig_2 = 0.026*	*sig_1 = 0.000* *diff = .120* *sig_2 = 0.000*	*sig_1 = 0.000* *diff = 33.778* *sig_2 = 0.000*	**Accepted**
7	**Effects of piracy**	piracy_rate	Spearman's rho	*N=196/740*	*sig = 0.000* *ρ = 0.350*	*sig = 0.000* *ρ = 0.453*	*sig = 0.000* *ρ = 0.391*	*sig = 0.000* *ρ = 0.342*	**Accepted**
8	**Effects of bandwidth**	int_bpp	Spearman's rho	*N=195/386*	*sig = 0.006* *ρ = -0.198*	*sig = 0.000* *ρ = -0.276*	*sig = 0.000* *ρ = -0.232*	*sig = 0.421*	**Rejected** (opposite holds)
9	**Effects of user education**	educ_ix	Spearman's rho	*N=196/547*	*sig = 0.000* *ρ = -0.354*	*sig = 0.000* *ρ = -0.301*	*sig = 0.000* *ρ = -0.381*	*sig = 0.000* *ρ = -0.261*	**Accepted**

5.3 Multivariate Regression Analysis

In the previous section we tested our hypothesis individually. The end result of such work is a list of associations that hold between the dependent variable(s) and some of the independent variables - some as expected, and some not. The next step is to move towards testing all the variables in our model *together*. This would enable us to assess the relative contribution of each of the variables; i.e., does a relation still hold when we control for the other factors? Multivariate regression analysis is a tool often used for this purpose.

An additional goal of regression analysis is to see what percentage of the sample variance can be explained using the variables that have been considered in the model. In this section we will perform regression analysis in two steps. The first step will be to use a simple multiple regression model. In the second step we will use *interaction terms* and *variable transformations* to increase the explanatory power of our model.

It should be stated that arguments against using regression analysis also exist. Most important of all, it might be argued that due to the limited number of independent variables at hand, our results will be very premature for explaining a socio-technical phenomenon as complex as botnets. Add to this all the proxies used, and the situation looks bleaker. Nevertheless, considering that a statistical model that explains botnets has not yet been proposed in the literature, this exercise, however partial, can act as a stepping stone towards more elaborate models.

5.3.1 The simple regression model

As stated, our first step will be to construct a linear multiple regression model – that is, without variable transformations, interaction terms, or higher order terms). This will be done by selecting variables, performing *stepwise regression*, and finally, checking the regression conditions.

Choice of variables

Table 28 lists one more time all the variables in the dataset, and indicates which ones will be included in the stepwise regression.

Table 28 – List of variables in dataset with notes on whether they will included in the regression model

Variable	Obs.	Associated with dep. variables?	Will be included in model?	Other notes
total_sub	741	yes, negative	yes	
market_share	709	no	no – total_sub which is better a proxy of size theoretically will be used[1]	
rev_persub	194	no	no - 'n' << obs, and will shrink dataset	
srv_cable	665	yes (only with *src_per*)	yes	dummy: 0, 1
lap_mem	741	yes	yes	dummy: 0, 1
cyber_mem	741	yes	yes	recoded to 0 (non-member) and 1 (signed/ratified/enforced)
piracy_rate	740	yes, positive	yes	
int_bpp	386	yes, negative	no - high association with other country level variables (in addition to low 'N').[1]	
educ_ix	741	yes, negative	yes	2008 data not available, so reused 2007 values for 2008

[1] When *market_share* and *int_bpp* where included in the model, the stepwise regression command would remove them.

Adding year to the dataset

Before starting, note that instead of using a separate 2007 and pool dataset, we opted to add 'year' as a variable to the regression. The reasoning is to capture changes over the years in the aggressiveness of the botnet herders. However adding time is only one strategy to do so, which might in fact not be the best, but most pragmatic. (The better strategy would be to capture this trend in a more elegant way, e.g., using a separate independent variable outside of our dataset.)[1] We will explore this issue some more at the end of this section.

Correlation matrix

Presented below is the correlation matrix between all the independent variables. The matrix scatter plots of all the non categorical variables are shown in Figure 77.

```
. pwcorr total_sub market_share rev_persub srv_cable lap_mem cyber_mem
piracy_rate int_bpp educ_ix, sig

             | 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
             |   0.0000
             |
  rev_persub |  -0.1149   0.1907   1.0000
             |   0.1107   0.0091
             |
   srv_cable |  -0.0958  -0.2110  -0.2721   1.0000
             |   0.0134   0.0000   0.0002
             |
     lap_mem |   0.1519  -0.1900  -0.1078   0.0922   1.0000
             |   0.0000   0.0000   0.1347   0.0174
             |
   cyber_mem |  -0.0729  -0.0887   0.0188   0.0738   0.1498   1.0000
             |   0.0474   0.0182   0.7946   0.0572   0.0000
             |
 piracy_rate |   0.0849   0.1284   0.1344  -0.0942  -0.3513  -0.6066   1.0000
             |   0.0208   0.0006   0.0617   0.0152   0.0000   0.0000
             |
     int_bpp |  -0.0527  -0.0759  -0.1881   0.0046   0.2420   0.4757  -0.5176
             |   0.3015   0.1393   0.0441   0.9316   0.0000   0.0000   0.0000
             |
     educ_ix |  -0.0835  -0.0793  -0.2180   0.1317   0.3184   0.4939  -0.6068
             |   0.0230   0.0348   0.0023   0.0007   0.0000   0.0000   0.0000
             |

             |  int_bpp  educ_ix
-------------+------------------
     int_bpp |   1.0000
             |
             |
     educ_ix |   0.3476   1.0000
             |   0.0000
             |
```

[1] Another option would be to leave out time altogether, lowering the explanatory power of the model, but also foregoing the caveats.

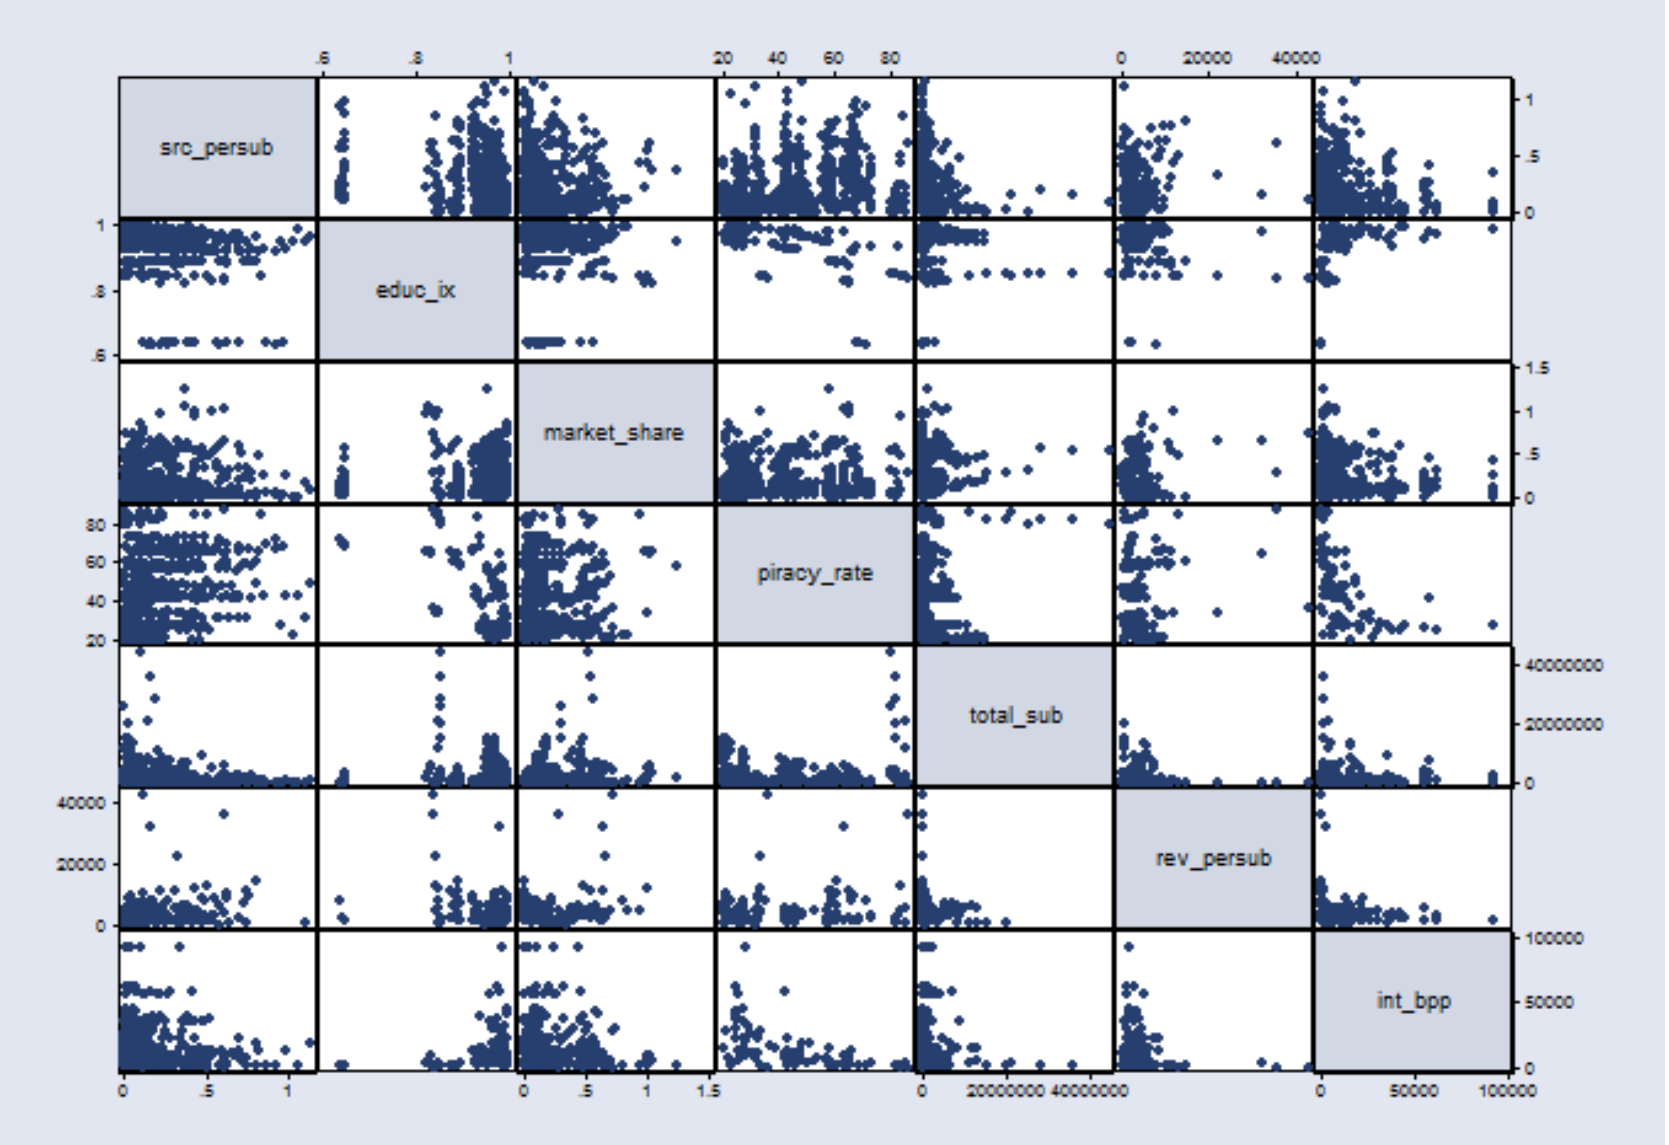

Figure 77 - Matrix plot of all non-categorical variables in dataset (2005-2008)

STEPWISE REGRESSION - USING SRC_PERSUB AS DEP. VARIABLE

The results of a stepwise regression with the chosen variables are as follows (backward selection is used, but forward selection yields the same output):

```
. sw reg src_persub cyber_mem educ_ix lap_mem piracy_rate srv_cable total_sub
year, pr(0.2)
                    begin with full model
p = 0.5931 >= 0.2000  removing educ_ix
p = 0.3132 >= 0.2000  removing cyber_mem
      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F(  5,    658) =   39.00
       Model |  6.00028781     5  1.20005756           Prob > F      =  0.0000
    Residual |  20.2464046   658  .030769612           R-squared     =  0.2286
-------------+------------------------------           Adj R-squared =  0.2227
       Total |  26.2466924   663  .039587771           Root MSE      =  .17541
------------------------------------------------------------------------------
  src_persub |      Coef.   Std. Err.      t    P>|t|                     Beta
-------------+----------------------------------------------------------------
   total_sub |  -8.34e-09   1.99e-09    -4.18   0.000                -.1480965
        year |   .0248756   .0061601     4.04   0.000                 .1391598
     lap_mem |  -.0534268   .0149817    -3.57   0.000                 -.132767
 piracy_rate |   .0040113   .0004301     9.33   0.000                 .3445915
   srv_cable |  -.0641315   .0147204    -4.36   0.000                -.1509358
       _cons |  -49.82627   12.36168    -4.03   0.000                         .
```

This is our basic regression model. The model is significant as a whole (F value is highly significant), all the betas are significant (t values are highly significant). The model explains **22%** of the variation. The sign of the betas are in the expected direction:

- negative for *total_sub*: bigger ISPs have lower levels of botnet activity
- negative for *srv_cable*: cable providers have lower levels of botnet activity
- positive for *year*: every year, the botnet phenomenon is getting worse
- negative for *lap_mem*: ISPs in countries that have signed LAP have lower botnet activity
- positive for *piracy*: ISPs in countries with higher piracy rates have higher botnet activity

The variables *educ_ix* and *cyber_mem* where dropped from the model.

Controlling regression assumptions

We will now control the regression assumptions. The results are presented in Figure 78. *Heteroscedasticity* is particularly observable in the *residuals vs. fitted values* plot. (We will leave remediation to the advanced model).The Q-Q plot to test the normality of the residuals is shown in Figure 79. There is a moderate deviation from the normal distribution present. Finally, the *variance inflation factor* does not indicate any particular case of *multicollinearity* (all <10).

```
    Variable |       VIF       1/VIF
-------------+----------------------
     lap_mem |      1.18    0.845787
 piracy_rate |      1.16    0.858603
   total_sub |      1.07    0.934305
   srv_cable |      1.02    0.976709
        year |      1.01    0.987179
-------------+----------------------
    Mean VIF |      1.09
```

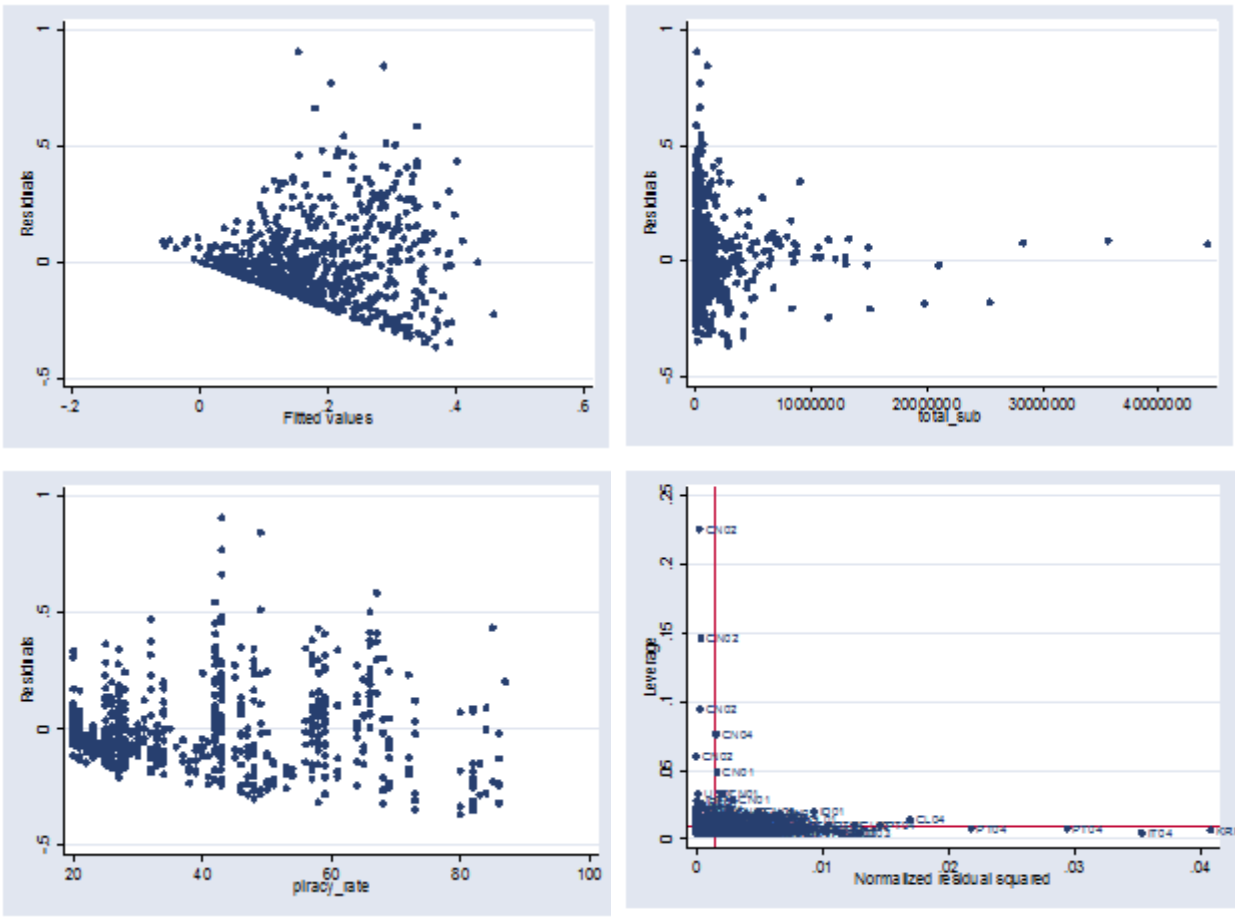

Figure 78 - Residual plots for linear regression model (src_persub as dep. variable)

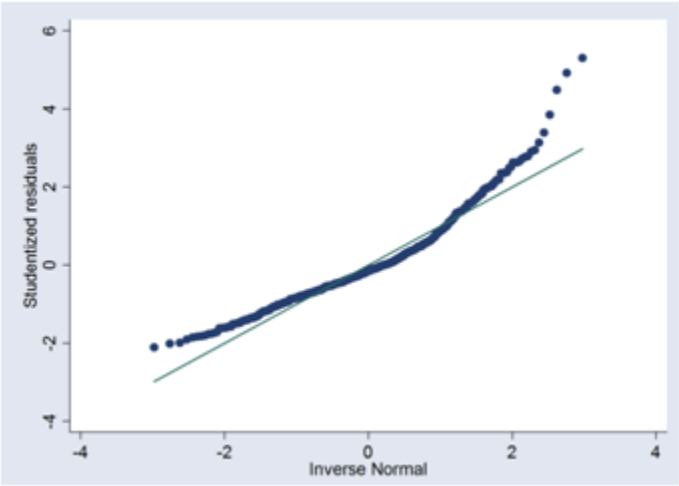

Figure 79 - QQ plot of residuals for linear regression model (src_persub as dep.)

STEPWISE REGRESSION - USING SPAM_PERSUB AS DEP. VARIABLE

The same model for *spam_persrc* is as follows:

```
. sw reg spam_persub  cyber_mem educ_ix lap_mem piracy_rate srv_cable total_sub
year, pr(0.2)  b
                     begin with full model
p = 0.7623 >= 0.2000  removing srv cable
p = 0.4012 >= 0.2000  removing cyber_mem
      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F(  5,   658) =   66.60
       Model |  1139090.46     5  227818.091           Prob > F      =  0.0000
    Residual |  2250692.86   658  3420.50586           R-squared     =  0.3360
-------------+------------------------------           Adj R-squared =  0.3310
       Total |  3389783.31   663  5112.79534           Root MSE      =  58.485
------------------------------------------------------------------------------
 spam_persub |      Coef.   Std. Err.      t    P>|t|                     Beta
-------------+----------------------------------------------------------------
   total_sub |  -3.44e-06   6.64e-07    -5.18   0.000                -.1702931
     educ_ix |   154.6848   47.82379     3.23   0.001                 .1324908
     lap_mem |  -12.86439   5.02912    -2.56   0.011                -.0889553
 piracy_rate |   1.588295   .173475     9.16   0.000                 .3796661
        year |   29.84135   2.05404    14.53   0.000                 .4645251
       _cons |  -60022.59   4122.966   -14.56   0.000                        .
```

The results are somewhat similar to the previous model. This model is significant (F is significant), all the betas are also significant (t-values are significant), and most of the signs are as expected (the one that is not is explained in the next paragraph). The model explains **33%** of the variance - higher than the previous model.

Two differences exist: First, *srv_cable* is removed from the model. We have seen something similar in the hypothesis testing: whether or not botnet activity levels are lower in cable providers depends on the metric used; it was not observed when *spam_persub* is used as the metric.

The second difference is not explainable for us: *educ_ix* is put back into the model, and with a direction reverse to that the individual hypothesis reveals.[1]

Controlling regression assumptions

The residual plots are shown in Figure 80. Heteroscedasticity is similarly particularly observable in the residuals versus fitted values plot. The normality plots looks similar to for src_persub. The VIF mean is 1.33 and does not indicate multicollinearity.

[1] If we manually take out educ_ix, then R^2_{adj} becomes 32.2%, which is still quite good. Thus one solution could be to just take out educ_ix.

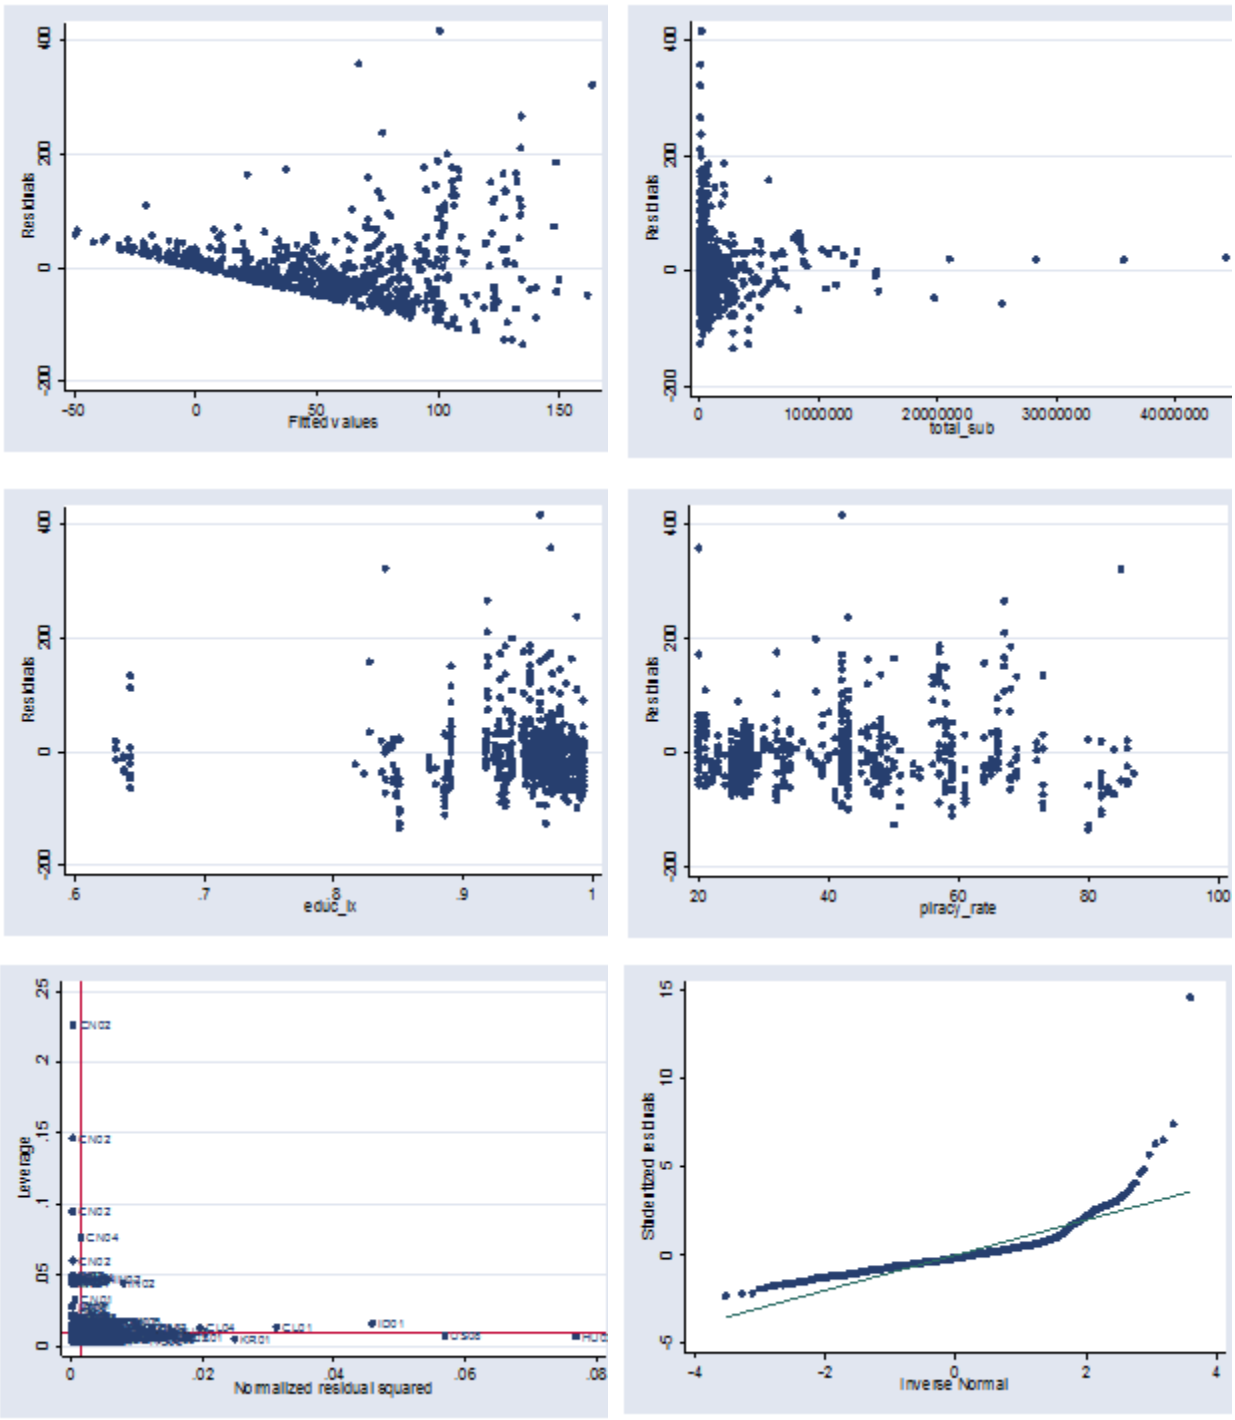

Figure 80 - Residual plots and QQ plot, for linear regression model (with dep. var spam_persub)

5.3.2 ADDING INTERACTION TERMS TO THE MODEL

In this section we will test the effects of various changes to our simple linear models, and seek to increase the explanatory power of the model (as measured by R^2_{adj}).

The first of these changes will be to add *interaction terms* between the variables *lap_mem*, *educ_ix*, and *piracy_rate.* The intuition is that these terms, which all reflect aspects of the environment the ISP operates in, influence each other, and are associated with each other (examples were given previously in section 5.2). Interaction terms, which are multiplication of such terms, are a common way to handle this scenario in regression analysis.[1]

STEPWISE REGRESSION – USING SRC_PERSUB AS DEP. VARIABLE

The output of the stepwise regression is as follows. We have added the four terms *lapXedu, lapXpir, eduXpir,* and *lapXeduXpir*. As can be seen, all the new terms have significant t-values, the overall model remains significant (F=0.0000), and R^2_{adj} increases to **33%**, from the previous 22%. Our hunch is correct. The only problem is that *total_sub* is now removed from the model – a variable which we expect to remain in the model.

```
. sw reg  src_persub  year total_sub srv_cable  cyber_mem lap_mem piracy_rate
educ_ix  lapXedu lapXpir eduXpir lapXeduXpir , pr(0.2)
                    begin with full model
p = 0.3389 >= 0.2000  removing total_sub

      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F( 10,   653) =   33.86
       Model |  8.96297064    10  .896297064           Prob > F      =  0.0000
    Residual |  17.2837217   653   .02646818           R-squared     =  0.3415
-------------+------------------------------           Adj R-squared =  0.3314
       Total |  26.2466924   663  .039587771           Root MSE      =  .16269

------------------------------------------------------------------------------
  src_persub |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
        year |    .021797   .0056969     3.83   0.000     .0106105    .0329835
 lapXeduXpir |   .1091919   .0195806     5.58   0.000     .0707433    .1476405
  srv_cable |  -.0791812   .0137599    -5.75   0.000    -.1062001   -.0521622
  cyber_mem |  -.0343001   .0192634    -1.78   0.075    -.0721259    .0035256
    lap_mem |   3.480345   1.177329     2.96   0.003     1.168538    5.792153
 piracy_rate |  -.0377809   .0124623    -3.03   0.003     -.062252   -.0133099
     educ_ix |  -2.878206   .8755132    -3.29   0.001    -4.597367   -1.159046
     lapXedu |  -3.865495   1.226949    -3.15   0.002    -6.274737   -1.456253
     lapXpir |  -.0980962   .0184121    -5.33   0.000    -.1342502   -.0619421
     eduXpir |   .0427058   .0130727     3.27   0.001     .0170362    .0683754
       _cons |  -40.83073   11.47599    -3.56   0.000    -63.36502   -18.29644
```

Two solutions come to mind for remediating the removal of *total_sub*. The first is to perform a log transformation on this variable, and the second is to add an interaction term with *srv_cable.* The idea for the first solution comes from the output of the *gladder* command in Stata (Figure 81). This command shows all the possible transformations of a variable, and as can be seen, the 'log' function transforms it to near normal. On the subject of transformations, and based on the same command, we also adopt a square root transformation on the depenent variable (see Figure 81). [2]

[1] Interaction terms are often used in variables related to demographics and institutional effects, as in these cases, typically, 'configurations' of variables makes more sense than the variable separately.

[2] What these transformations actually mean will be discussed later in section 5.3.4

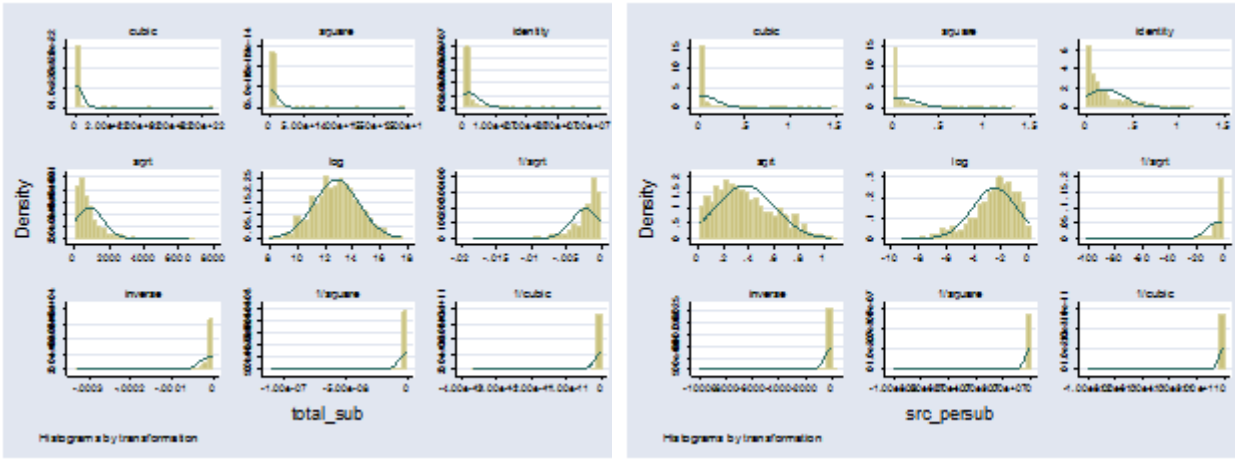

Figure 81 - Histogram of all possible transformations of the variable total_sub(left) and src_persub(right)

The intution behind the second solution comes from the fat that our explanation for higher security in ISPs that provide cable access, or are larger, were similar; i.e., the factors work in the same direction. Our explanation was that these ISPs most probably use automated botnet detection and mitigation methods – the first group due to having the infrastructure already in place, and the second group due to the necessacity to have securiy that scales. Thus we would expect the presence of both factors at the same time to be less influential then each of them alone, and an interaction term can be used to capture this. (An example is given in 5.3.4)

Finally, *cyber_mem* is removed as its beta wasn't significant. The result of this new regression analysis is as follows. The result is an increase in R^2_{adj} to **36%**. The F value of the model is highly significant, and so are all the t-values. Using *src_persub* as the dependent variable, this regression model as almost the best we can get.

```
. sw reg   src_per_sq   year totsub_ln srv_cable   cblXsubln lap_mem piracy_rate
educ_ix  lapXedu lapXpir eduXpir lapXeduXpir , pr(0.2)
                    begin with full model
p < 0.2000          for all terms in model

      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F( 11,    652) =   35.41
       Model |  12.4336775    11  1.13033432           Prob > F      =  0.0000
    Residual |  20.8116487   652  .031919707          R-squared     =  0.3740
-------------+------------------------------           Adj R-squared =  0.3634
       Total |  33.2453262   663   .05014378          Root MSE      =  .17866

------------------------------------------------------------------------------
  src_per_sq |      Coef.a   Std. Err.      t    P>|t|                    Beta
-------------+----------------------------------------------------------------
        year |    .0187167   .0063501     2.95   0.003                .0930338
   totsub_ln |   -.0385053   .0123797    -3.11   0.002               -.1226287
   srv_cable |   -.4617647   .1274027    -3.62   0.000               -.9656368
   cblXsubln |    .0670622   .0228187     2.94   0.003                .782244
     lap mem |    5.093856   1.293589     3.94   0.000               11.24734
 piracy_rate |   -.0344505    .013446    -2.56   0.011               -2.629586
     educ_ix |   -2.844884   .9559523    -2.98   0.003               -.7780779
     lapXedu |   -5.650225   1.348427    -4.19   0.000               -12.02671
     lapXpir |   -.1265761   .0203935    -6.21   0.000               -12.08912
     eduXpir |    .0387899   .0141324     2.74   0.006                2.463882
 lapXeduXpir |    .1411135   .0216976     6.50   0.000                12.3424
       _cons |   -34.26849   12.76148    -2.69   0.007                      .
```

Controlling regression assumptions

The residual plots are presented in Figure 82. We have not included the residual plots for the interaction terms as that would not be meaningful. High heteroscedasticity can still be seen in the residuals versus fitted values plot. (Please note that a log transformation of src_persub, instead of a square-root one, removed the heteroscedasticity but didn't increase R^2). The Q-Q plot shows an increased normality. The VIF output is also not meaningful when interaction terms are used - as they obviously have multicollinearity – and hence, not given.

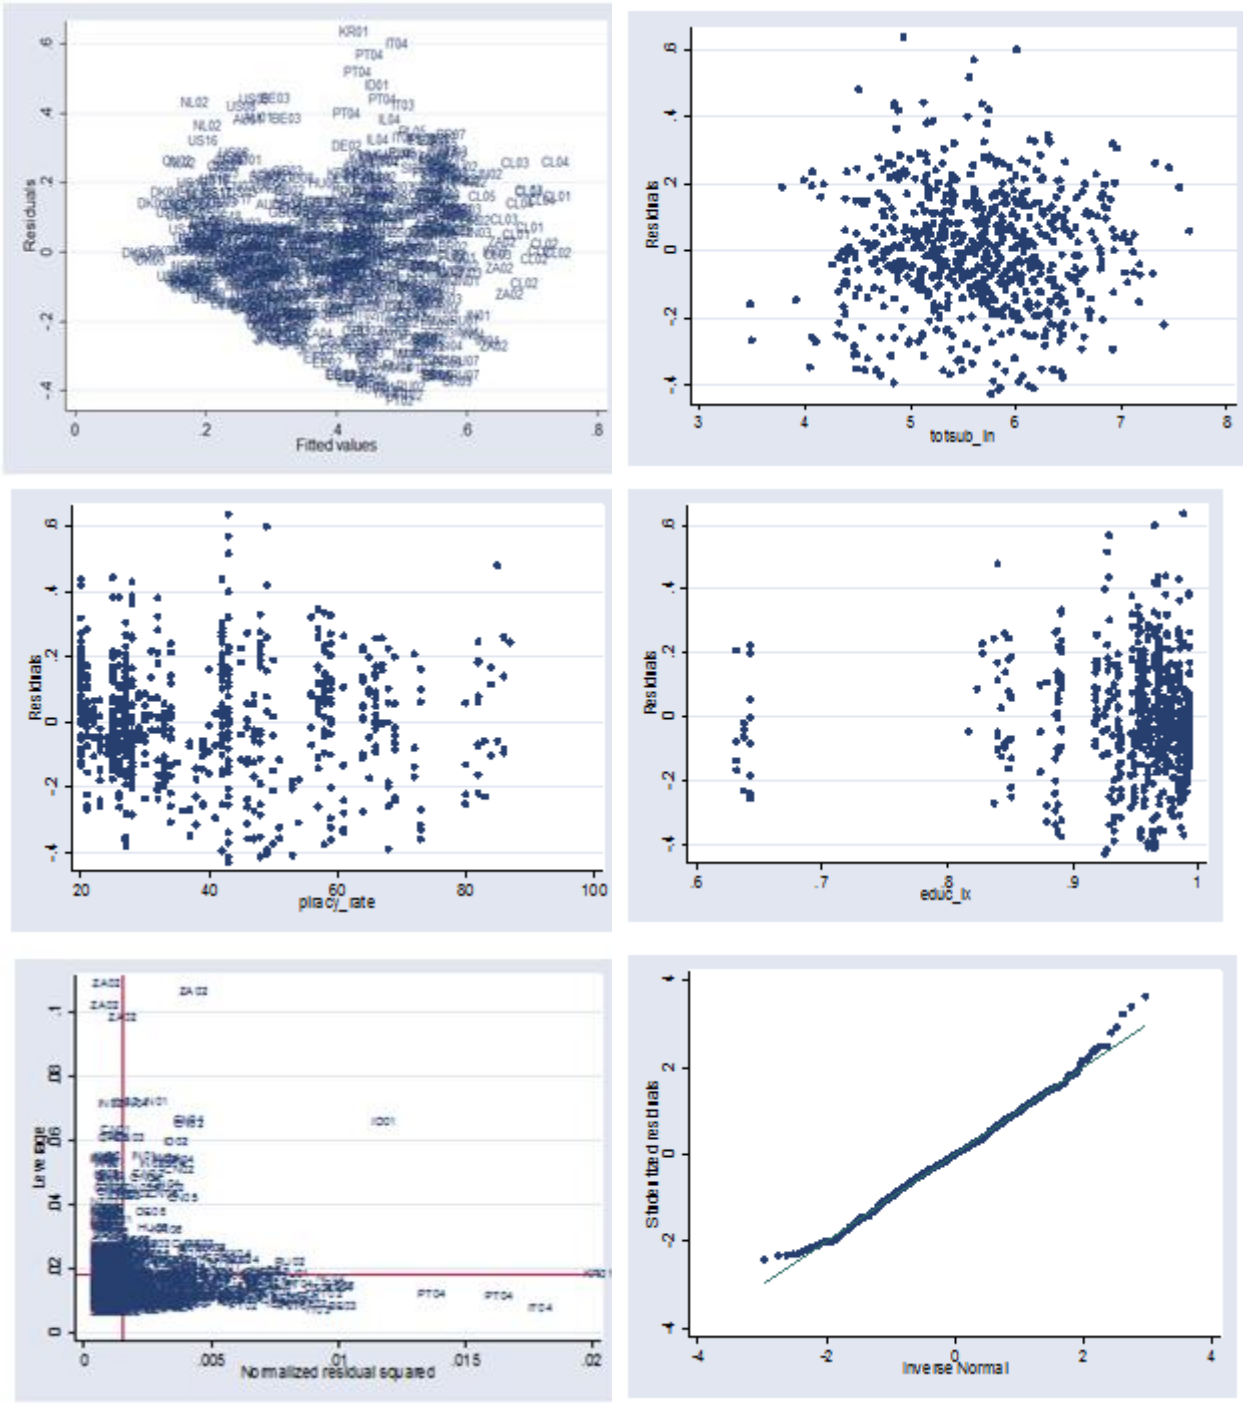

Figure 82 - Residual plots and QQ plot for the interaction regression model (src_persub as dep. var.)

STEPWISE REGRESSION – USING SPAM_PERSUB AS DEP. VARIABLE

Based on the experience of the previous heading, and as a *confirmative step*, we will straight away use the model with the interaction terms, and chosen transformations. The result is given below. As can be seen, R^2_{adj} astonishingly rises to **46%** - from the previous 31%. The model remains significant (both F-value, and t-values).

```
. sw reg   spam_per_sq   year totsub_ln srv_cable   cblXsubln cyber_mem lap_mem
piracy_rate educ_ix  lapXedu lapXpir eduXpir lapXeduXpir , pr(0.2)
                    begin with full model
p = 0.7820 >= 0.2000  removing cyber_mem
      Source |       SS       df       MS              Number of obs =     664
-------------+------------------------------           F( 11,   652) =   54.13
       Model |  5293.71898    11   481.24718           Prob > F      =  0.0000
    Residual |  5796.68362   652  8.89061905           R-squared     =  0.4773
-------------+------------------------------           Adj R-squared =  0.4685
       Total |  11090.4026   663  16.7276057           Root MSE      =  2.9817
------------------------------------------------------------------------------
  spam_per_sq |      Coef.   Std. Err.      t    P>|t|                    Beta
-------------+----------------------------------------------------------------
        year |   1.877936   .1059781    17.72   0.000                .5110739
   totsub_ln |  -1.099565   .2066085    -5.32   0.000               -.1917276
   srv_cable |  -5.201773   2.126254    -2.45   0.015               -.5955745
   cblXsubln |   .9279622   .3808265     2.44   0.015                .5926339
 lapXeduXpir |   1.462732   .3621157     4.04   0.000                7.004666
     lap_mem |   44.01195   21.58902     2.04   0.042                5.320659
 piracy rate |   -.541715   .2244043    -2.41   0.016               -2.263883
     educ_ix |  -39.03307   15.95412    -2.45   0.015               -.5844978
     lapXedu |  -48.02652   22.50422    -2.13   0.033                -5.59698
     lapXpir |  -1.345696   .3403519    -3.95   0.000               -7.036904
     eduXpir |   .6705663   .2358588     2.84   0.005                2.332035
       _cons |  -3721.893   212.9795   -17.48   0.000                       .
```

Controlling regression assumptions

The plots for controlling the regression assumptions are given in Figure 83. The plots are rather similar to what we saw for the src_persub model, although the outliers have changed. (Interpretation of the residuals requires further work; it might aid in making inferences regarding the nature of the variables missing from the model.)

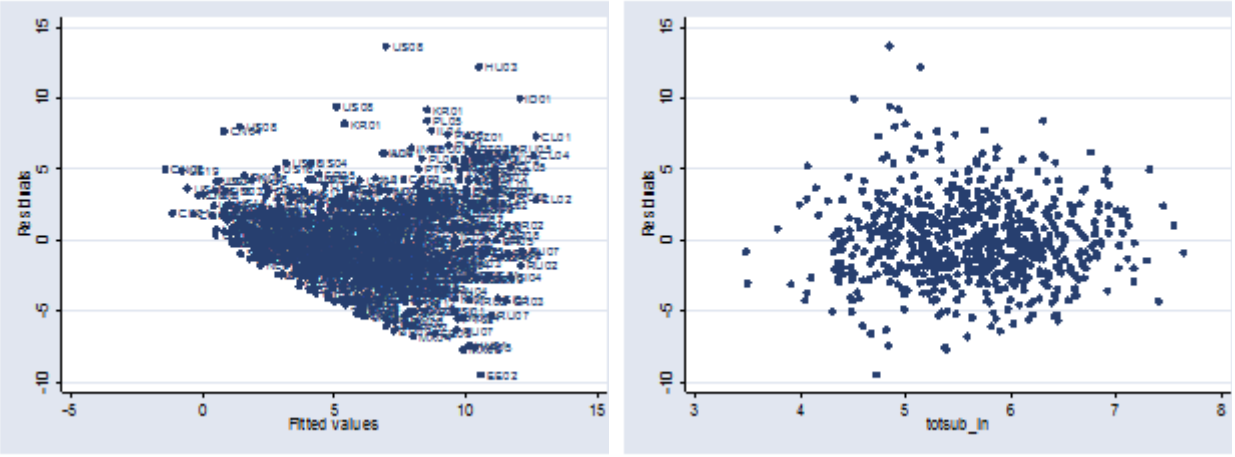

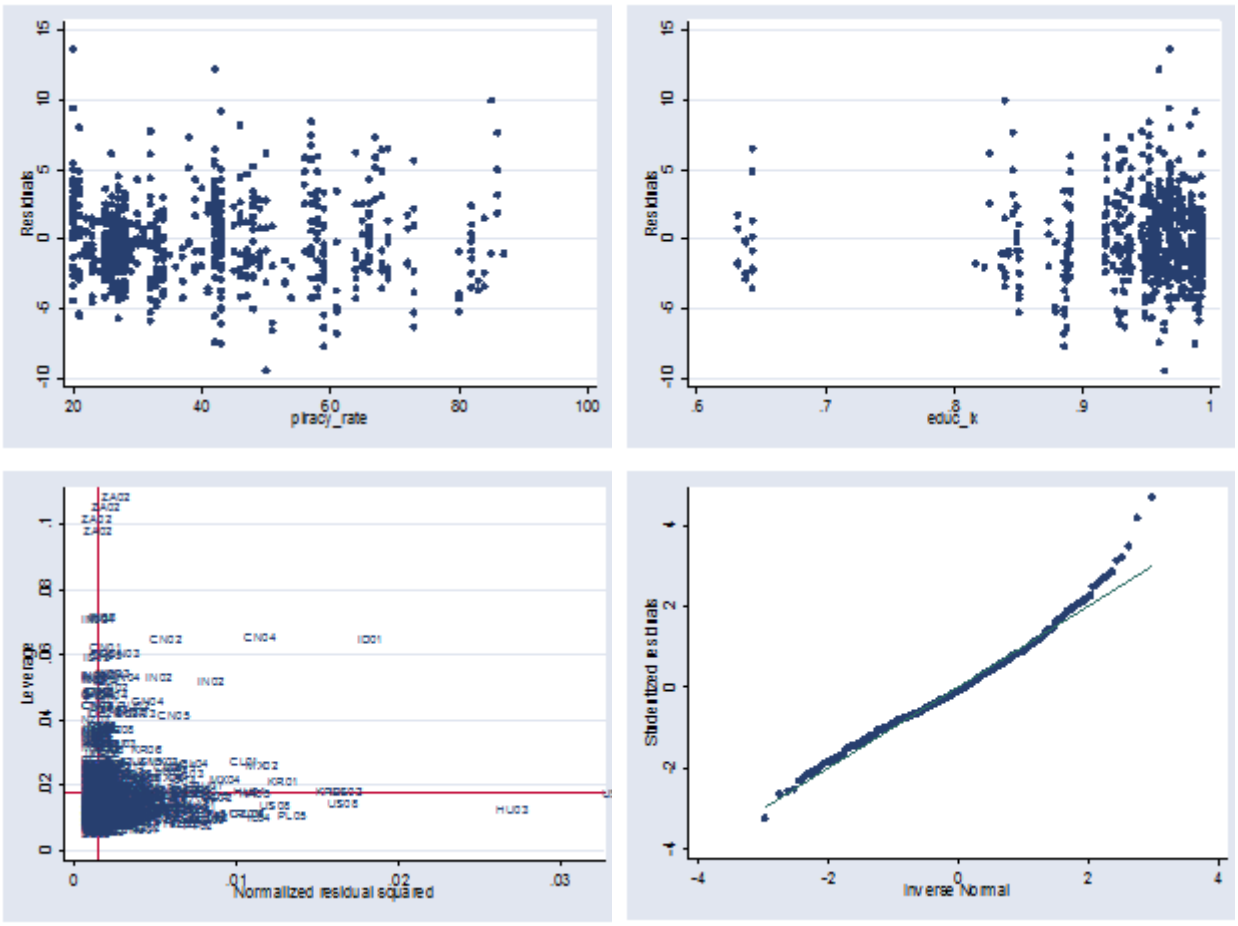

Figure 83 - Residual plots and QQ plot for the interaction regression model (spam_persub as dep. var.)

Experimenting with the regression model

Some other experiments with the regression terms were done to see if the predictability of the model (R^2) would increase, without success. These are listed in Table 29 for the curious reader.

Table 29 - Experiments with the final regression model

Procedure	Effect on R^2_{adj}
Removing outliers (e.g., *src_persub* > 0.92).	R^2 increased only slightly – decided to keep outliers in the model.
Using factor analysis to find the latent variable behind *educ_ix, piracy_rate*, and *lap_mem,* and entering that term in the model.	Regressions actually had a lower explanatory power.
Creating a regulatory index, *reg_ix = cyber_mem + lap_mem* , and using that in place of *lap_mem* and *cyber_mem* in the model.	Made absolutely no difference.
Adding *cyber__mem* as a fourth interaction term to *lap_mem*piracy_rate*educ_ix*.	Increased R^2 only slightly - not worth the complexity it would add to the model.
Other transformations (higher order terms, and transformations of the other variables)	No significant improvement.

5.3.3 Model interpretation

Explaining the addition of the year variable

Adding time to a regression model implies the capturing of an effect (i.e., part of the sample variance) that is not otherwise represented in the model. The rationale for the decision to add year to our model was the dynamism of the botnet phenomenon - in particular, the increased aggressiveness of the botnet herders over the years. This is not the ideal strategy, but rather the most pragmatic (– a better strategy would be to capture this trend using new independent variables from outside of our dataset).

Additionally, adding year as a direct variable, instead of a dummy, can also raise concerns. To investigate these issues, we decide to conduct a series of tests on different ways the year variable could be used (or not used) in the model, and compare the end results. These tests were as follows:

- Removing the year variable altogether
- Adding the year variable in the form of a dummy
- Running the regression on individual years, separately

The last step also doubles as ***model confirmation***: a step usually undertaken after constructing a regression model in which the dataset is partitioned in various ways, to see if the model holds in the subsets as well.

Table 30 shows the results. As can be seen, the results are inconclusive - different variables are dropped in different years, and the R^{2}'s change in a considerable range. This illustrates more than anything else the changing nature of the botnet phenomenon, and suggests that the years do not have the same dynamics.

Table 30 - Testing the final regression model on subsets of the data

	src_persub as dependent variable	spam_persub as dependent variable
Pooled (2005-2008) data – (current model)	*F = 0.000* *R^2_{adj} = 0.363*	*F = 0.000* *R^2_{adj} = 0.468*
Pooled (2005-2008) data, without the year variable	*F = 0.000* *R^2_{adj} = 0.356*	*F = 0.000* *R^2_{adj} = 0.214*
Only 2005 data	*F = 0.000* *R^2_{adj} = 0.388 (piracy dropped)*	*F = 0.000* *R^2_{adj} = 0.253 (educ_ix, piracy dropped)*
Only 2006 data	*F = 0.000* *R^2_{adj} = 0.3925 (none dropped)*	*F = 0.000* *R^2_{adj} = 0.343 (educ_ix, piracy dropped)*
Only 2007 data	*F = 0.000* *R^2_{adj} = 0.330 (totsub dropped)*	*F = 0.000* *R^2_{adj} = 0.343 (lap, cable dropped)*
Only 2008 data	*F = 0.000* *R^2_{adj} = 0.357 (totsub, piracy dropped)*	*F = 0.000* *R^2_{adj} = 0.399 (none dropped)*
Year entered as a dummy variable	*F = 0.000* ***R^2_{adj} = 0.376***	*F = 0.000* ***R^2_{adj} = 0.483*** *(dum_yr06 insignificant)*

INTERPRETATION OF THE MODEL TERMS

We finish this section by interpreting the final regression models we created. This is as follows:

$$\sqrt{src_{sub}} = 0.02 \times year - 0.04 \times \ln tot_{sub} - 0.46 \times srv_{cable} + 0.07 \times cblXsubln + 5.1 \times lap_{mem} - 0.03 \times piracy_{rate} - 2.8 \times educ_{ix} - 5.7 \times lapXedu - 0.04 \times lapXpir + 0.04 \times eduXpir + 0.14 \times lapXeduXpir - 34 + \varepsilon$$

$$\sqrt{spam_{sub}} = 1.9 \times year - 1.1 \times \ln tot_{sub} - 5.2 \times srv_{cable} + 0.93 \times cblXsubln + 44 \times lap_{mem} - 0.54 \times piracy_{rate} - 39 \times educ_{ix} - 48 \times lapXedu - 1.3 \times lapXpir + 0.67 \times eduXpir + 1.5 \times lapXeduXpir - 3722 + \varepsilon$$

Coefficient of determination

The coefficient of determination is approximately 36% when *src_persub* is used as the dependent variable and 47% when *spam_persub* is used. These values would not normally be considered high, but taking into account the fact that botnets are a complex socio-technical phenomenon, it is impressive that such percentages of the variance among ISPs can be explainable, with the help of just a few variables.

Interpretation of the transformations

Two of the terms use transformations. From a practical point of view, these transformations were chosen because they normalized the underlying variable. However, we have theoretical explanations to back these choices:

- ***Logarithmic*** transformation is used when the ***'order of a magnitude of a variable'*** is more important than its absolute number. This is completely the case with number of subscribers, which we are using as a proxy for ISP size. For us, an ISP that has 5.4 million users and one with 5 million users is the same (i.e., they would have similar policies, equipment, etc, if the only determinant was size). However, an ISP with 50,000 users and 450,000 users would differ considerably. Hence, the order of magnitude is what matters.

- ***Square root*** transformation is used when the ***'law of diminishing returns'*** is in effect. This makes sense, as *src_persub*, which is a proxy for percentage of an ISP's users that are infected, will never reach 100% - some users will just never be infected, irrelevant of the ISP security policies. It will also never reach 0% - some users will always be infected, whatever the policy. Hence, the law of decreasing marginal returns holds.

Interpretation of the betas

Presented below are the interpretations of the coefficients of both the interacting and non-interacting terms in the model:

- ***Year:*** every year, a 2% increase the infection rate of the population occurs. This fits with the literature – botnets and malware are on the rise. (Do bear in mind the explained limitations on interpreting the time variable in the model)

- ***Convention on Cybercrime***: the fact that the variable *cyber_mem* is removed from the regression, while *lap_mem* is not, is rather interesting. It basically means that the London Action Plan is more important and effective when it comes down to botnets. The reason could be that the Convention on Cybercrime deals more with criminal law in general, where as the LAP, is more of a regulatory activity, and is focused specifically on spam and malware.

- ***The interaction terms between the ISP level variables (total_sub and srv_cable):*** The idea behind adding this interaction term is that since both these factors work in the same direction (i.e., we hypothesize that

the larger ISPs, and cable providers, both have an infrastructure in place that eases botnet mitigation for them), the presence of both of these factors together will be less influential than expected, e.g., a big ISP probably already has automated botnet mitigation, whether it provides cable or not, so if it's also providing cable wouldn't change this issue much.

The coefficients (betas) confirm this idea: *ln_totsub* and *srv_cable* both have negative coefficients, e.g, the number of bot infections drops by 4% everytime the size of the ISP doubles. Their interaction term however has a positive sign, meaning that if both an ISP is large and a cable provider, the drops are offset by a certain amount.

- ***The interaction between the country level variables (educ_ix, lap_mem, piracy_rate):*** The idea behind this interaction term is that a) together they constitute the operating environment of the ISP, and b) they are highly associated with each other, e.g., countries with higher education index are more likely to have lower piracy rates, and have signed the London Action Plan. In other words, to a certain extent these variables will be measuring similar things, and changing together *(i.e., they find meaning in certain configurations – often the case with demographic and institutional variables)*. This necessities the addition of interaction terms to boost (or decrease) their combined effects.

  Interpretation of the sign of the coefficients in the country level interaction terms requires further work, and can be done by entering actual values for these terms in the equation, and seeing their net change (also known as *calculating the elasticity*).

The regression model depicted

In previous chapters, the factors explaining the security performance of ISPs in botnet mitigation has been visualized in the form of a *theoretical model* and a *measurement model.* Based on the results of this chapter, we can visualize this relation also in the form of a regression model presented in Figure 84.

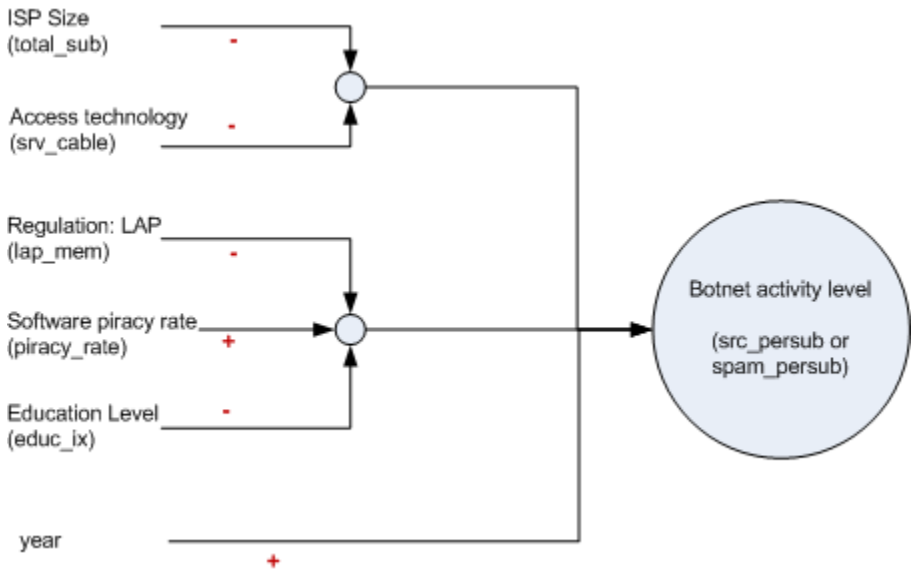

Figure 84 - Regression model depicted

CHAPTER 6 – CONCLUSIONS

6.1 ANSWERING THE RESEARCH QUESTION

We have already come a long way in answering our research question! In this concluding chapter, we will revisit our research question, and based on the empirical evidence (i.e., results of the data analysis), see how they can be answered. The main research question was as follows:

> *"Are Internet Service Providers crucial intermediaries in botnet mitigation efforts? Do they significantly differ in the degree in which they mitigate botnets? If so, to what extent can these differences be explained? And what implications does this have for policy?"*

Based on the empirical findings, we can without difficulty answer the first three parts of this question. The last part is however a prescriptive question, and requires some reflection on the findings. [1] This section is structured as follows: first we will review the empirical results one by one, and contemplate what each hypothesis test result means for policy making. We will then provide a summary, and present some final recommendations. These recommendations will without doubt be subject to much debate, which will be followed in the subsequent 'discussion' section.

6.1.1 REVIEWING THE FINDINGS

You will recall that our final regression model had an approximate explanatory power of 40%. While this can be considered high, it also indicates missing variables in our model. Since these missing variables could change the meaning of the results, a word of caution is raised several times during the interpretation of the findings. Also please note

H1: ISPs are central

The acceptance of this hypothesis points to the existence of the 80/20 rule in the realm of botnet activity, i.e., 80% of the global botnet activity originates from less than 20% of all Autonomous Systems. This is good news for governments – the bulk of the botnet problem is concentrated within a small number of economic players, so governments can tackle the problem by getting into talks with only a handful of organizations in their countries.

H2: ISPs differ significantly

This hypothesis was also accepted. It was found that first of all, the increase in the number of people going online is the main driver for the rise of botnets. This means that experts shouldn't be extremely alarmed by the statistics pointing to the worsening of the botnet phenomenon. Despite this, it was also found that there are orders of magnitude differences between retail brand ISPs, when corrected for size, in regards to the level of botnet activity occurring on their networks. This leaves plenty of room for improvements, and for negotiations by the policy makers with the lacking parties.

H6: Effects of regulation[2]

[1] Regarding the number of sub-questions, note that in sections 1.2 and 2.4 we identified five while here we list four. The missing sub-question, "what are botnets and why is important to mitigate them", was answered in the literature review.

[2] The hypotheses are reviewed in the order that makes the narration more logical, not their original numbering.

The acceptance of this hypothesis shows that regulation seems to be effective in combating botnets. This is again good news for governments, as it shows that regulation in the area of botnets can fix some of the market failures that are occurring. Before contemplating further though, we must raise caution that this finding might simply be the by-product of another cause, e.g. some common characteristic of the country that influences both people's online behaviour and the decision by policy makers to pass such laws – the so called *correlation does not imply causation fallacy*.

Interestingly enough, specifically targeted regulation, such as those stimulated by the London Action Plan, seems to be more effective in mitigating botnets than broader ones, such as the Convention on Cybercrime. Several reasons can be thought of on why this happens:

- One is that the LAP is related to regulatory activity, which is closer to the market than criminal prosecution, and hence more effective.
- Another is that the effects of more specific regulation might actually be in resolving the legal ambiguity that ISPs face in tackling botnets (e.g., regarding privacy issues), and not related with the threat of prosecution.
- Yet another explanation could be that regulatory activity in the area of spam and malware, is a proxy of a government and private sector intent on resolving these issues, i.e., it is an indicator, not the actual cause.

H3: Effects of size

Contrary to what is stated in the literature, bigger ISPs actually perform better in terms of good net citizenship (i.e., emitted junk), when corrected for the size of their subscriber base. This finding has several implications.

The most likely reason that comes to mind for this finding is that large retail ISPs, due to their scale, deploy automated tools to detect network attacks, respond to abuse notifications, and quarantine infected users.[1] More empirical data would be needed to test these claims, and if they do turn out to be true, then governments should encourage the adoption of automated botnet mitigation technologies by ISPs.

Another implication is that the consolidation of telecoms does NOT have any security implications at this point in time (– something that might be expected to happen as their market power increases). This could be due to the fact that the amount of competition is already quite high in the telecom market (with 14 tier ISPs and over 200 major players in just the OECD countries).

A third and more speculative implication could be that in the incentive structure of large retail ISPs, reputation (i.e., brand image) plays a bigger role than peer pressure. If this turns out to be true, then public name and shame campaigns on players with bad security would be very effective.

H4: Cable providers versus DSL providers

It was found that cable providers have higher security then DSL providers, when measured in terms of number of bot infections on their networks. We can think of two reasons why this turns out to be the case, both with outreaching implications.

Our original hypothesis was that cable providers have already in place an infrastructure for monitoring and controlling network traffic - necessitated by the fact that bandwidth is shared in their infrastructure. If this turns out to be the true cause of this observation, then (similar to the automation argument,) governments should encourage all ISPs to install such equipment, and upgrade their infrastructures as necessary. This encouragement

[1] In the same line of reasoning, they are also more likely to have better network equipment in place.

can be done through various mechanisms, such as giving loans, negotiations, or even by regulation. Additionally, governments should fund research and innovation into monitoring and control technologies that tackle botnets.

An alternative cause that comes to mind is that since cable providers typically have more retail users then DSL providers, it is easier for them to impose restrictive network policies (e.g., policies such as blocking port 25 traffic, which business users would be opposed to). If this turns out to be true - that stricter network policies can be effective, as long as there is no user backlash involved - then regulation can be helpful. If the law mandates ISPs to impose such policies on all their users, then there will be no user backlash, or competitive reason not to do so. (This argument is similar to the legal ambiguity argument given in H6).

H7: Effects of software piracy

We accepted the hypothesis that software piracy rates are associated with increased botnet activity levels. One explanation often given is that either the owners of pirated software typically do not update their software or install security patches, thus making them more vulnerable to malware. (This is because either the software refuses to be updated, or the user fears that the update will break the crack.) If this turns to be the case, then governments should mandate software vendors to release security patches for all users, even those using illegal copies. A totally different solution could be to run campaigns against software piracy.

Other reasons are also cited, [1] but what is important is to watch out for the *"correlation does not imply causation fallacy"* here. Software companies have incentives in painting software piracy as the root of all evil, where as the truth might be that, as an example, countries with lower piracy rates are also the same ones that have adopted the LAP, and actually, the adoption of the LAP is the reason for lower botnet activity in such countries.

H9: Effects of education

Although this hypothesis was accepted, and it was found that higher education levels associate with less botnet activity, we feel uncomfortable in basing any policy implications on this result. The reasons is that although we do expect increased online security awareness of users to equate to less bot infections, our finding is most probably a by-product of some other cause, as education index is very crude indicator for online awareness.[2]

H8: Effects of bandwidth

The literature mentions broadband penetration to be one of the technology enablers of malware and botnets. Our bandwidth hypothesis is, contrary to this believe, empirically rejected - we have actually found bandwidth and botnet activity to be *negatively* associated.

This could have several reasons. One could be that bandwidth is a proxy of how well an ISP is run, in addition to how good an infrastructure it has, i.e., well run ISPs manage to provide higher bandwidths and also mitigate botnet activities. Another reason could be that what we are observing is in fact a be a by-product of certain other factors changing in chorus with increased access speeds at the country level, i.e., countries with high broadband penetration are more likely to adopt the LAP, etc. (Adding the average access speeds of the ISPs themselves instead of the countries they reside in, would help explore this issue.)

Despite the uncertainties in explaining the rejection of this hypothesis, it still contains good news for governments: increasing broadband penetration does NOT automatically equal less security, as the effects of broadband as an enabler for botnets can be easily offset by other factors.

[1] e.g., that the use of peer-to-peer networks, often a source of malware, is higher among users of pirated software. (We cannot verify this claim.)

[2] This suspicion is increased by the fact that in some of the stepwise regressions, this variable *educ_ix* was dropped.

SUMMARY

We have summarized the major points from the above discussion in Table 31. The first three parts of the research question are answered in the box below.

> ***Taking a look back at the research question, we can state that: yes, ISPs are crucial intermediaries in botnet mitigation efforts; and yes, they differ significantly in the degree in which they mitigate botnets. Among the factors investigated, the most promising are that targeted regulation seems to be effective; and the fact that larger retail ISPs and cable providers have lower botnet activity levels.***

Table 31 - Summary of the major empirical findings

Based on	Finding
Tests of individual hypothesis	ISPs are the focal point in botnet mitigation: approximately 200 retail ISPs account for 80% of the bot infections in these countries; in other words, the bulk of the problem is concentrated within a small number of economic players.
	Retail brand ISPs differ significantly in regards to the level of relative botnet activity occurring on their networks. Variability among ISPs of similar size suggests that ISPs do in fact face different tradeoffs, and more importantly, their choice of security practices makes a big impact.
	The number of subscribers is negatively associated with botnet activity levels. This contradicts the commonly held belief that larger ISPs perform worse in terms of security.
	Cable providers have a better security performance then DSL providers – on average 10% lower bot infections. (Speculative reasons for this could be existence of traffic monitoring systems due to the shared bandwidth infrastructure; or stricter network policies due to more residential users).
	Targeted regulation such as those stimulated by the LAP appear to be effective - ISPs operating in countries that have signed the 'London Action Plan' have on average 13% lower bot infections. Conversely, the broader 'Convention on Cybercrime' appears to be ineffective.
	Other results include finding piracy rate to be positively associated with botnet activity levels; and ARPU and market share to not seem to influence have any significant relation;
Multivariate regression analysis	Approximately 40% of the sample variance regarding the relative degree in which ISPs mitigate botnets (i.e., number of infected sources corrected for size), can be explained using the variables subscriber count, cable access, LAP membership, privacy rate, education index, and year. (This percentage is high, considering the limited number of variable used in explaining a complex phenomenon)
	The interaction terms among the country level variables (LAP membership, piracy rate, and education index) in the regression model indicate that these variables move in configurations – often the case with demographic and institutional variables.
	An interaction term also exists between cable access and subscriber count. The direction of the beta strengthens the speculation that the increased security performance in these organizations has a common cause, most possibly the use of automated abuse monitoring and handling (~anti-bot) systems.

6.1.2 POLICY IMPLICATIONS OF THE EMPIRICAL FINDINGS

By combining the major empirical findings, we can draft two broad policy recommendations for governments, as follows:

Recommendation 1: policy makers should engage in dialogue with the ISPs regarding the mitigation of botnets

The rationale behind this recommendation is pretty straightforward: the compelling evidence that ISPs are the focal point in botnet mitigation. The contents of the discussion should be regarding the *scope* on which ISPs need to be taking action, and on the mechanisms that work best. A good starting point would be to compare the security metrics of the ISPs with their peers. As we know that ISPs differ significantly, this comparison can give transparency to the scale of the efforts ISPs are currently undertaking. (It can additionally show the laggers that room for improvement exists; and provide a basis for comparing the effectiveness of security policies.) The dialogue with ISPs will hopefully shed light on many of the issues at stake.

Recommendation 2: policy makers should encourage the adoption of technologies that automate and enable large-scale bot identification, remediation, and abuse handling

This recommendation is built on the speculation that the lower bot infection rates of large ISPs and cable providers is indeed due to the use of use of automated abuse monitoring and handling systems. [1] To recap, this idea stems from the following facts:

- *Large ISPs perform better in terms of security performance*. One possible explanation is that in order to scale, these ISPs have had to deploy automated tools for handling tasks such as detection of network attacks, responding to abuse notifications, and quarantining infected users.
- *Cable providers perform better then DSL providers.* One explanation (and our original proposition) is that due to their 'shared bandwidth infrastructure', these providers had to put in place systems for monitoring and controlling network traffic, which can also be used for bot detection.
- In the regression model, an interaction term exists between subscriber count and cable access, with the sign of the coefficient opposite to that of the individual terms. This suggests that the effects caused by being large and those caused by being a cable provider are similar, as the existence of both of them has less of an effect than the summing their individual effects.[2] The "automated tools effect" would fit this explanation.

[1] It should be mentioned that the use of automated abuse monitoring and handling tools is not the standard – many ISPs are known to only handle such tasks manually, for a variety of reasons. For instance, experts from the ISP XS4ALL told us that their customers actually *expect* to be notified of abuse problems in *person*. The result is that XS4ALL doesn't fully automate its abuse handling process, but rather puts many FTEs on the task.

[2] This could be due to the boring reason of interdependency between two variables but gladly, they are independent.

As you can see, the common explanation is that the availability of technologies that automate bot mitigation tasks plays a significant role in the reduction of botnet activity levels. The consequence of accepting this idea is: encourage the adoption of such technologies.[1]

Please note that both of these policy recommendations are based on interpretations of *our* data, in light of *our* research question. But this is not the complete picture. For instance, using only our data, we cannot say anything about the possible side effects of these recommendations, should they be implemented. We will reflect extensively on this matter in section 6.2. Only then will we have a satisfactory understanding of the policy implications.

Reviewing the ENISA recommendations in light of our work

In a report commissioned by the European Network and Information Security Agency (ENISA), Anderson and his colleagues (2008b) discuss some of the practical options available to EU governments for improving security failures. Their report draws out, from both economic principles and empirical data, a set of fifteen recommendations about what information security issues should be handled at the Member State level and what issues may require harmonisation – or at least coordination by the states.[2] Here we evaluate some of their recommendations in light of the empirical evidence we have gathered in this project.

Recommendation 3: ENISA collect and publish data about the quantity of spam and other bad traffic emitted by European ISPs

We agree with this recommendation. This research (and the broader OECD project it is part of) shows that collecting spam data is valuable, and such data can be used by academics to generate fresh insights into the botnet phenomenon, and it can also open room for negotiations between regulatories and ISPs - based on actual facts.

Recommendation 4: The EU introduce a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines...

We agree with part of this recommendation that ISPs are central to the botnet problem, and that governments should engage with them in order to mitigate the botnet problem. We are however not sure if fining ISPs is the best path to pursue, or whether more encouraging methods should be used.

Recommendation 7: Security patches be offered for free, and patches be kept separate from feature updates

In light of the acceptance of the *hypothesis on piracy*, we agree with this recommendation.

Other statements in the report

The report contains several other statements that can be verified using our data. These are presented in the following table.

Table 32 - Verification of several claims in the ENISA report

Claim	Agree or disagree?
There is great variation among ISPs ...	Agree (H2)
Bigger ISPs are worse...	Disagree (H3)
Network consolidation is a barrier to security	Disagree, at least in regards to telecoms (H3)
Robust metrics to measure security should be developed	Agree – see recommendations section for several ideas

[1] If we take this idea one step further, research and development in such technologies should also be encouraged (by grants, government purchases, etc). However, if adoption of these technologies on a large scale occurs, the market might itself take the role of encouraging innovation in this area.

[2] The list of fifteen proposals was presented in Chapter 2, Table 9.

6.2 DISCUSSION

Throughout this thesis, we have argued that ISPs are the focal point in the fight against botnets, and towards the end, we provided empirical evidence to support this proposition. Such a recommendation, although shared by many security experts, is far from controversial, and many valid concerns have been raised against it. These concerns are not so much as to whether shifting responsibility to ISPs would be effective in mitigating botnets, but are rather related to the bigger picture: things that could possibly break and go wrong if we adopt such an initiative; and answering questions such as whether this is a sustainable solution, or deciding who has to pay for the costs involved. In this section we will review some of these arguments. We will first review reasons why ISPs should take the leading role in the fight against botnets, then present the major arguments against this proposition, and finally, try to reconcile these opposing viewpoints.

6.2.1 REASONS FOR ISPS TO BE MORE ACTIVE

Arguments in favour of ISPs taking a more active stance boil down to two things: practicality and efficiency. Before listing any reason, we must be reminded that the costs of inaction against botnets, or maintaining the status quo, are simply too high. Botnets are massive criminal enterprises used not just for spamming, but also for identity theft, financial theft, DDoS attacks, and many other not-so-friendly things. The problem of malware and botnets is growing, and the damages (due to both fraud and the productivity loss caused by the malware infections) are in the orders of billions of Euro[1]. Current defences are struggling.

Botnets and malware, manifestations of cyber-insecurity, are complex socio-technical problems and the results of many failures and causes: intent cyber-criminals, naive end-users, poorly written software, insufficient law enforcement, path dependency, etc. Tackling most of these issues, although necessary, is not at all easy. The number of actors involved is huge, for instance consider the approximately billion people online, a portion of which will always remain unaware and unprotected from the risk of malware infections.

This is where Internet Service Providers become attractive: as intermediaries providing online access to the bulk of Internet users, they are in a unique position to mitigate the botnet problem: they have access to the end-users; they can act as gatekeepers between their computers and the broader Internet; they are technically more competent than users to manage these risks; and there are far fewer ISPs then end-users. ISPs can often detect that a subscriber's computer has been compromised, due to the unusual traffic patterns of that subscriber – for instance, abnormally high email activity is most likely spam from a botnet. They can then quarantine the infected subscriber, and redirect her to a website with instructions for remediation.[2] Our data provides compelling evidence for stating that ISPs are the focal point: approximately 200 ISPs account for 80% of the bot infections in the enlarged-OECDs countries[3] (- making it practical for policy makers to get into discussions with). From an economic standpoint, making ISPs responsible for bot-infections of their customers is also efficient: responsibility is assigned to a party that is most capable of managing the risks (due to their expertise and position).

For precisely these reasons, many security experts have argued for increasing the role of ISPs in the fight against botnets: Eric Davis, head of Google's Anti-Malvertising (Naraine 2009); Anderson and his colleagues in the ENISA report (2008b); Livingood and his colleagues for the IETF (2009); Dave Rand, CTO of Trend Micro (independent.ie 2008); and many others (see for example IIA 2009; MAAWG 2009; Economist 2009; Wash 2007). Our data further

[1] Please refer to section 2.1.2 for actual statistics.

[2] The number of measures that ISPs can perform is quite extensive, falling in nine broad categories – see section 2.3.3.

[3] These countries together have approximately 375 million broadband subscribers and 1.2 billion Internet users.

shows great variation among the retail ISPs in terms of botnet activity levels on their networks – indicating room for manoeuvre and improvement among many of them.

ISPs are already taking measures, but in a mostly reactive manner (Van Eeten and Bauer 2008). If ISPs do not take on a more proactive stance against botnets, (and unless some other practical mitigation solutions magically come to being), far reaching negative consequences might be encountered. Zittrain (2009) gives one example: *"...intentional inaction at the network level may be self-defeating, because consumers may demand 'locked-down' endpoint environments that promise security and stability with minimum user upkeep".* When endpoints are locked down, malware is kept out, but unfortunately, developers will also be unable to deliver their innovative products directly to users.

6.2.2 Reasons against increasing the role of ISPs

The arguments against ISPs taking a more direct and vigilante role in botnet mitigation can roughly be grouped as follows: ineffectiveness of this initiative in the long run; creating room for opportunistic behaviour by different parties; negative net-effects on society; and fairness.

Ineffectiveness in the long run

Some expert believe that benefits gained by assigning responsibilities such as identifying, notifying, and remediating bots to ISPs will only be short lived - despite all the efforts that will go into implementing them.

One argument given is that detection of bots is technically *unreliable.* Detection of bots is similar to finding meaning in a series of bits (network traffic). This is not easy, as although currently 'known' bots exhibit traffic patterns that make many of them identifiable, it is quite possible that newer bots will get around the identification systems, especially considering the adaptive behaviour of the botnet herders. Detection is further complicated by the risk of 'false positives' – traffic that appears to be from a bot, but in fact isn't. Quarantining users due to false positives will result in some extremely unhappy customers. Thus, ISPs will err on the side of caution, making the detection system even less effective.[1]

Another argument is that even if botnets were totally mitigated, the overall situation of cyber-insecurity and cyber-crime wound not sustainably improve, as the underlying causes of the botnet phenomena (such as profitability of cybercrime, immunity from prosecution, naivety of users, and buggy software) still remain. The only effect will be the criminals changing tactics and moving on to the next profitable malicious activity; the mitigation efforts will simply become a stage in the on-going 'arms race'.[2][3] An example more sustainable solution would be to go after the key payment channels associated with online crime (Anderson 2007).

We discussed these issues with Dave Rand recently, and questioned whether shifting responsibility on to the ISPs would make a lasting difference? He gave us an interesting answer: these issues might (and most probably will) occur, but this fact does not disqualify the endeavour, as in his view, the ultimate goal is to force the criminals to constantly innovate and evolve, and by doing so, raise their transaction costs. Otherwise, the scale of the attacks would become much larger.

[1] One could think of using external metrics to assess the performance of ISPs and as a tool to push them towards more action; such metrics will however in many cases be contested by the ISPs - "we're doing fine and these metrics aren't a good indicator".

[2] Snowshoe spamming was one example mentioned in this thesis of post-botnet spamming techniques.

[3] In the short run, the implementation of the quarantining system could itself create new abuse opportunities! An example would be showing fake notification pages to subscribers in order to sell them scare-ware and rogue AV software!

POSSIBILITY OF OPPORTUNISTIC BEHAVIOUR

Room for opportunistic behaviour will be created when ISPs are given the responsibility to mitigate botnets. For one, such a responsibility necessities the implementation of systems for identifying bot infected subscribers. ISPs might be tempted to use such a system as an opportunity to classify user traffic for their own revenue-generating applications (i.e., targeted advertisement). The responsibility also authorizes ISPs to disconnect infected users. In extreme cases, they might use their newly acquired powers to disconnect users or traffic that they do not wish to carry, e.g., by labelling it malicious or risky. (This concern touches on a current debate regarding the *network neutrality* principle[1].)

Another party that might want to use the monitoring capabilities of ISPs to promote their own agenda are the copyright owners. For years, they have been trying to stop the sharing of copyrighted material on the net, often asking ISPs to filter out content. ISPs have often responded that this is not technically possible (as an example see OUT-LAW 2008). The copyright owners will see the newly acquired capabilities of ISPs as an opportunity to push for identification and blocking of copyright infringing content. Such a move would not only be unfavourable with users, but would also create headaches and unasked responsibilities for ISPs.

Other possibilities of opportunism include the creation of a *'moral hazard'* for end-users, and the exhibition of *'free rider'* behaviour by software vendors. Believing that their ISPs will be protecting them, some users might decrease their security vigilance and adopt risky online behaviour (e.g., stop using anti-virus software; clicking on what-ever comes along; etc). Software vendors might become lax in providing timely patches, instead demanding that ISPs block certain worms.

The final group we will look at are the politicians. They might be tempted to frame this initiative as the magical cure for all cyber-insecurity problems, and forego or stall some of the other steps that need to be taken in this regards (such as stepping up law enforcement). In the end, some of these opportunistic behaviours might be mitigated by supplementary regulation, oversight, transparency, etc. In practice however these additional measures might turn out to be too tricky or costly to implement.

NEGATIVE NET-EFFECTS ON SOCIETY

Even in the absence of opportunistic behaviour, implementing anti-bot measures can have large scale negative impacts that might well out-weight the benefits..

Foremost, implementing the various anti-bot measures (such as quarantining users) are all costly endeavours.[2] [3] Who will pay for them? Due to the scope of the phenomenon, ISPs will most likely be unable to cope with these costs alone. One outcome could be the introduction of a 'security surcharge' – i.e., increasing the subscription rates for ISP access, so that security becomes part of the fee users pay. This increase in Internet subscription fees might result in the slowing down of broadband adoption, thus creating opportunity costs for society as a whole.

[1] *'Net-neutrality'* is a principle for access networks participating in the Internet to impose no restrictions on content, sites, platforms, or modes of communications. The proponents of this principle fear that else, telecoms might use their infrastructure to block content of their competitors, or create artificial scarcity to oblige their subscribers to pay more, etc (Wikipedia 2010b).

[2] The costs constitute of the cost of implementing the feature itself (e.g., required software, devices, and man-power); and the added customer support costs required to provide guidance to the disconnected customers who contact call support. This second portion of the cost is actually the more expensive part (Van Eeten and Bauer 2008).

[3] Additionally, increased transaction costs due to the possibility of law-suits filed by users who have been mistakenly cut-off from the net are another cost to reckon.(The users might have been in critical situation, such as an outright emergency, or a student who has to hand in a report before a deadline. Even if they do not sue the ISP, they and society have incurred costs).

Another outcome could be that governments decide to subsidize the efforts. This effectively spreads the costs to tax-payers as a whole, which from a social welfare perspective might turn out to have a negative net outcome.

Non-monetary risks also exist. One is that certain fundamental rights of the people, such as the right to privacy[1], or the right to freedom of speech, might suffer - as a result of the ISPs been allowed to "snoop" upon all user traffic by default (i.e., without needing a warrant first). These issues become more serious if *deep packet inspection* technology is employed.

Additionally, transferring responsibility of infected systems to ISPs will likely result in stricter network controls and restrictions for end-users. This would make it harder for the "layman" to run his own mail-server or web-server. Putting aside statements such as "this was not how the Internet was meant to be", imposing network restrictions would still be undesirable, as they might create a *trajectory* that in the long-term slows down Internet innovation, restricts *hacktivism*[2], and has other negative social outcomes.

Are these risks and foreseeable costs, worth the risks of damages that bot infected computers would inflict, should they remain online?

FAIRNESS

A final criterion to consider is the concept of fairness. An important question that can be asked is whether it is 'fair' to ask ISPs to take responsibility for a problem that they have not caused, simply because they can? They are after all, 'mere conduits' according to current law.

Another question in this realm is whether remediation services will be provided for users of all operating systems, or only for those using Windows? It will surely be unfair to users of other OSes if the services are provided only for Windows users, especially if the costs are being paid by all subscribers.

6.2.3 RECONCILIATION

How can we reconcile the arguments for and against ISPs taking the lead in fighting off botnets? On the one hand, botnets are one of the (if not *the*) most serious threats that the Internet economy faces; the costs of inaction against them are high; and among the available options, for the reasons mentioned, increasing the role of ISPs seems to be the most promising solution. On the other hand, there are good arguments against such an approach - questioning the long-term effectiveness, the possibility of opportunistic behaviour, the net effects on social welfare, etc. The reality is that we are faced with a political decision, and balancing these arguments is something that needs to be done by the policy makers.

If the decision is made for ISPs to take a more active role, then solutions must be adopted to mitigate the expressed risks. As an exercise, we have thought of several example solutions:[3]

- Seeing the proposal (of increasing the role of ISPs) as part of a larger package that includes complementary initiatives such as user education, and increased law enforcement efforts.

[1] e.g., against unreasonable searches

[2] Hacktivism is "the nonviolent use of illegal or legally ambiguous digital tools in pursuit of political ends. These tools include web site defacements, redirects, denial-of-service attacks, information theft, web site parodies, virtual sit-ins, virtual sabotage, and software development." (Samuel 2007) The goal is to produce similar results to those produced by regular activism or civil disobedience, such as promoting free speech and human rights.

[3] Please note that the suggestions are only examples - so they may be controversial and have "rough" edges at this point.

- Limiting the use of deep packet inspection technology, and instead advocating less intrusive detection techniques - such as monitoring outbound email volume, or access to known C&C domains.
- Providing the opportunity for users to opt-out of the service (should they be willing to accept the risks).
- Separating the entity that decides what constitutes botnet activity (and provides the identification steps) from the entity that enforces the rules (i.e., the ISPs).
- Spreading the costs of the initiative among multiple actors. As an example, these actors could include the ISPs, software vendors, broadband users, and the criminals themselves (via fines collected from them!).
- Involving all stakeholders in discussions regarding the implementation of the initiative, and being fully transparent about the system that is implemented, the risks it creates, and the checks and controls used to mitigate these risks.

To reiterate, these are only examples, and the actual decisions are a matter for policy makers to consider, taking into account all the costs and benefits for society.[1] Like the solution to other complex social phenomenon, increasing the role of ISPs will likely be an iterative process with much learning and refinement involved. Luckily, there are already some innovative initiatives underway.

One initiative recently proposed in Germany, is for the government to set up an advisory centre that will help users purge their computers of viruses and bots (H-Online 2009). In this way, the biggest cost in quarantining infected machines is absorbed by an entity other then the ISPs, changing the incentive structure. In other news, in the U.S., the mega ISP Comcast has began trials of "the Constant Guard security program". One feature of this program, called "Service Notice", lets customers know whether they are infected with a bot (Comcast 2009).[2] Follow up studies on the outcomes of these initiatives (and similar efforts) is highly recommended, as it will help answer the broader questions surrounding botnet mitigation and cyber-insecurity.

[1] A good question would still be: how best to incentivise ISPs, should the call be made?

[2] Incidentally, Comcast has been involved in quite a number of controversies regarding customer support and network neutrality over the past few years (Wikipedia 2010a)!

6.3 Limitations and Suggestions for Further Research

No research is complete without a discussion on the reliability, validity and limitations of the work, and a list of suggestions for further research. These issues will be investigated in the closing section of this thesis.

Before delving in to the issue of limitation, a distinction must be made between two forms of limitations. The usual definition of limitation is regarding the execution of the research (i.e., the reliability of the measurement tools and the validity of the methods used in answering our research question). But another form of limitation also exists, which has more to do with the overall scope of the project: is the research question that we have answered, relevant to solving the broader, underlying cyber-security problem? These two forms are usually phrased as *"did we answer the question right?"* versus *"did we answer the right question?"*. We will investigate the answer to the first question formally. The second question is more a matter of reflection and was discussed in section 6.2.

Reliability of the instruments

Reliability has to do with repeatability of the measurements (Velde, Jansen, and Anderson 2007); since the empirical data for this research has come from secondary sources, the issue of repeatability does not hold for us - we will end up with the same dataset no matter how many times we rerun the scripts. The issue of reliability could be raised for the spam trap data itself, but since the spam trap has collected its data at a certain time in the past, repeatability will again be a non-issue – although the validity of the spam trap as a sample does need to be examined.

Validity of the instruments and the research strategy

Validity is the extent to which an instrument measures what we hope it measures; and the extent to which a research strategy results in the type of conclusions that we draw from it (Velde, Jansen, and Anderson 2007). Much can be said regarding these two questions.

We examined the validity of our measurement instrument in depth in the Methodology Chapter, including questions of whether spam is a valid proxy of botnet activity, etc. We have relisted the main points in Table 33. You should recall that that most of the limitations in the measurement instruments are common among similar types of research.

The validity of the results depends on a number of things, among them, the validity of the sample. This was examined in the Data Preparation Chapter, and the spam trap deemed to be a representative sample of world-wide spam traffic. Other questions include whether the correct statistical techniques were used in hypothesis testing, and whether correct conclusions are drawn from them. The major issue here, which was covered extensively in the Data Analysis Chapter, concerns the limited number of independent variables we have, and the fact that most of them are proxies for some other measurement that we seek (increasing the risk of mistaking association for causation). Table 33 relists these points as well.

Table 33 - Limitations of this research

Category	Description	Details
Validity of measurement of the dep. var.	Limitations in assuming spam as a proxy of botnet activity	Ch 3
	Limitations in assuming IP addresses (unq_srcs) as a proxy for infected subscribers (considering dynamic IPs, NAT, etc)	Ch 3
	Distortions of the spam_msgs metric	Ch3
	GeoIP inaccuracies	Ch 3
	Mistakes in operator to ASN mapping	Ch 4
Validity of measurement of indep. var.	Most of our independent variables come from secondary sources, such as the Worldbank or TeleGeography, asserting their validity, so no particular problem should exist here.	-

Category	Description	Details
Validity of results: sample validity	Is the spam-trap source a representative sample of world-wide spam traffic? (Seems so)	Ch 4
	TeleGeography contains only a subset of retail ISPs operating in each country. Their choice of ISPs might create a systematic bias in the selection of ISPs in our sample.	
Validity of results: content validity	The limited number of independent variables for explaining a complex phenomenon, some of which are crude proxies, can be problematic.	Ch 3 Ch 5
Validity of results: construct validity	Interactions between independent variables were not fully examined. The fallacy of 'correlation does not imply causation' thus may occur in certain cases	Ch5
	Analysis of actors with regards to botnets does not take into account actor interdependencies	Ch 2
	Conceptual framework does not incorporate dynamic effects	Ch 2

Suggestions for further research

This research work can be extended in several ways. One is to improve the measurement of the dependent variable (e.g., by adding more sources of malicious traffic, or building new metrics based on the current data); Another is to enrich the number of independent variables in the dataset (this can come from a variety of sources, such as directly surveying ISPs, or using other secondary sources); And last, but not least, performing post-qualitative research - interviews with various categories of ISPs - can add context and valuable insights to the findings. Of course, the scope could also be broadened to cover some of the fundamental issues raised in the discussion section. These suggestions are listed in Table 34.

Table 34 - Suggestions for further research

Category	Suggestions
Improving measurement of the dependent variable (botnet activity)	Adding other data-sources of malicious traffic (SANS, Conficker, etc)
	Building a more robust dependent variable, which would be a weighted combination of unq_src & spam_msg, and using this for statistical testing and regression analysis
	Developing the persistence metric
	Categorizing ISPs (vigilant, rogue, etc)
	Gathering country level information on ISP policies regarding Dynamic IPs and NAT
Enriching the battery of independent variables	Surveying ISPs directly regarding their security policies, market, environment, incentives, etc (this option would be probably exceedingly expensive)
	Getting such data from ENISA, which has performed similar surveys
	Adding blacklisting data (e.g., from Spamhaus)
	Adding data on privacy laws (-> legal ambiguity)
	Adding data on ISP age and R&D budget (-> effects of innovation)
	Add data on ISP users and market segment (e.g., % business to non business)
	Even more complex variables: • How end-user security behaviour can be modelled, so that the effect of ISP security decisions can be more clearly highlighted? • How ISP security culture can be modelled – this was brought up several times in the OECD/XS4ALL workshop as an important factor in security decisions
Data analysis	• Performing "panel data" tests on our data – might reveal additional findings • Regression model: calculating elasticity; investigating year; investigating residuals
Post qualitative study on the ISPs	Interviewing experts from ISPs based on different groupings, e.g.: • under-performers , normal, and over-performers • small and the large ISPs • cable and DSL providers in order to check some of the points raised in our policy implications, such as finding the security measures are really effective, etc.
Ideas regarding the broader context	Follow up studies on the success of the current trials underway in Germany, the U.S. (Comcast), and other countries.
	Investigating the various issues raised in the discussion regarding the arguments against such an initiative – such as reliability of bot detection, DPI, privacy implications, trajectories, etc.

REFERENCES

Anderson, R. 2001. Why information security is hard - an economic perspective.

———. 2007. Closing the Phishing Hole – Fraud, Risk and Nonbanks.

———. 2008. *Security Engineering: A guide to building dependable distributed systems*: Wiley.

Anderson, R., R. Böhme, R. Clayton, and T. Moore. 2008a. Security Economics and European Policy. In *Workshop on the Economics of Information Security*.

———. 2008b. Security Economics and the Internal Market.

Anderson, R., and T. Moore. 2007. Information security economics - and beyond.

ArsTechnica. 2007. Deep packet inspection meets 'Net neutrality, CALEA. (13/10/2009), http://arstechnica.com/hardware/news/2007/07/Deep-packet-inspection-meets-net-neutrality.ars/4.

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.

BBC. 2007. Criminals 'may overwhelm the web'. (25 January 2007), http://news.bbc.co.uk/2/hi/business/6298641.stm.

———. 2008a. Net firms reject monitoring role. (15 February 2008), http://news.bbc.co.uk/2/hi/technology/7246403.stm.

———. 2008b. Study shows how spammers cash in. (10 November 2008), http://news.bbc.co.uk/2/hi/technology/7719281.stm.

———. 2009a. Botnet 'ensnares government PCs'. (21 April 2009), http://news.bbc.co.uk/2/hi/technology/8010729.stm.

———. 2009b. Microsoft bounty for worm creator. (13 February 2009), http://news.bbc.co.uk/2/hi/technology/7887577.stm.

———. 2009c. US launches cyber security plan. (29 May 2009), http://news.bbc.co.uk/2/hi/americas/8073654.stm.

———. 2009d. Zombie computers 'on the rise'. (19 October 2009), http://news.bbc.co.uk/2/hi/entertainment/8032886.stm.

BERR. 2009. Information Security Breaches Survey 2008. Department for Business, Enterprise & Regulatory Reform.

Blorge. 2007. Storm Worm network shrinks to about one-tenth of its former size. http://tech.blorge.com/Structure:%20/2007/10/21/2483/.

Cisco. 2007. Botnets: The New Threat Landscape. Cisco Systems.

———. 2008. Cisco 2008 Annual Security Report.

Comcast. 2009. Comcast Announces Constant Guard Program. http://www.dslreports.com/forum/r23152256-Comcast-Announces-Constant-Guard-Program#23152256.

CSI. 2008. CSI Computer Crime & Security Survey 2008.

Economist. 2009. Batten down the cyber-hatches. (30 April 2009), http://www.economist.com/world/europe/displaystory.cfm?story_id=13569241.

ENISA. 2007. Provider Security Measures: Survey on Security and Anti-Spam Measures of Electronic Communication Service Providers.

ETIS. 2007. Best Practices in Anti-SPAM: Advisory document for the ETIS community.

Goodman, J., G. V. Cormack, and D. Heckerman. 2007. Spam and the ongoing battle for the inbox. *Communications of the ACM* 50 (2):24–33.

GoogleTechTalks. 2007. An Economic Response To Unsolicited Communication. In *YouTube*.

H-Online. 2009. Germany to set up centre to coordinate fight against botnets http://www.h-online.com/security/news/item/Germany-to-set-up-centre-to-coordinate-fight-against-botnets-880077.html.

Hammer, M. 2009. A Few Thoughts on the Future of Email Authentication. *CircleID*, http://www.circleid.com/posts/20090414_thoughts_on_future_of_email_authentication.

IBM. 2008. IBM Internet Security Systems X-Force 2007 Trend Statistics.

———. 2009. IBM Internet Security Systems X-Force 2008 Trend & Risk Report.

IIA. 2009. Internet Service Providers Voluntary Code of Practice for Industry Self-Regulation in the Area of e-Security.

independent.ie. 2008. Setting the security trend. http://www.independent.ie/business/technology/setting-the-security-trend-1391313.html.

IndustryCanada. 2005. The Digital Economy in Canada - Appendix B. Industry Canada, Government of Canada

IronPort. 2008a. Internet Malware Trends 2008 - Storm and the future of social engineering.

———. 2008b. Internet Security Trends 2008.

KasperskyLab. 2008. Kaspersky Security Bulletin 2007: Spam report.

Kerbs, B. 2009. Security Fix - Verizon to Implement Spam Blocking Measures http://voices.washingtonpost.com/securityfix/2009/02/verizon_to_implement_spam_bloc.html.

Koppenjan, J., and J. Groenewegen. 2005. Institutional Design for Complex Technological Systems.

Livingood, J., N. Mody, M. O'Reirdan, and Comcast. 2009. Recommendations for the Remediation of Bots in ISP Networks. *IETF*.

Loder, T., M. Van Alstyne, and R. Wash. 2004. An economic answer to unsolicited communication.

MAAWG. 2005. Managing Port 25 for Residential or Dynamic IP Space - Benefits of Adoption and Risks of Inaction. Messaging Anti-Abuse Working Group.

———. 2007a. Best Practices for the Use of a Walled Garden. Messaging Anti-Abuse Working Group.

———. 2007b. BIAC and MAAWG Best Practices for Internet Service Providers and Network Operators. Messaging Anti-Abuse Working Group.

———. 2009. Common Best Practices for Mitigating Large Scale Bot Infections in Residential Networks. Messaging Anti-Abuse Working Group.

MAAWG, and APWG. 2006. Anti-Phishing Best Practices for ISPs and Mailbox Providers. Messaging Anti-Abuse Working Group and Anti-Phishing Working Group.

MAPS. 2009. *MAPS - Stopping Spam at its Source* 2005 [cited 25 October 2009]. Available from http://www.mail-abuse.com/.

McAfee. 2009. The Carbon Footprint of Email Spam Report. McAfee.

MessageLabs. 2008. MessageLabs Intelligence: 2008 Annual Security Report. Symantec.

———. 2009. MessageLabs Intelligence: Q2/June 2009 Symantec.

Microsoft. 2009. Microsoft Security Intelligence Report (Volume 6, July through December 2008).

Mueller, M. 2009. Cyber-security for people? Or nations? *IGP Blog*, http://blog.internetgovernance.org/blog/_archives/2009/1/31/4076192.html.

NAE. 2008. Grand Challenges for Engineering: National Academy of Engineering http://www.engineeringchallenges.org/Object.File/Master/11/574/Grand%20Challenges%20final%20book.pdf.

Naraine, R. 2009. Google exec calls for ISPs to get tough on botnets. http://blogs.zdnet.com/security/?p=4404.

NetEqualizer. 2009. What Is Deep Packet Inspection and Why the Controversy? « NetEqualizer News Blog. http://netequalizernews.com/2009/06/22/what-is-deep-packet-inspection-and-why-the-controversy/.

OECD. 2005. Anti-Spam Toolkit of Recommended Policies and Measures.

———. 2007. Malicious Software (Malware): A Security Threat to the Internet Economy. (DSTI/ICCP/REG(2007)5/FINAL).

———. *About OECD* 2009. Available from http://www.oecd.org/pages/0,3417,en_36734052_36734103_1_1_1_1_1,00.html.

OUT-LAW. 2008. Belgian ISP wins reprieve in copyright infringement filtering case. http://www.out-law.com/page-9537.

Provos, N., D. McNamee, P. Mavrommatis, K. Wang, and N. Modadugu. 2007. The ghost in the browser analysis of web-based malware.

Rittel, H. W. J., and M. M. Webber. 1973. Dilemmas in a general theory of planning. *Policy sciences* 4 (2):155–169.

Samuel, A. 2007. Overview of dissertation on Hacktivism and the Future of Political Participation. http://www.alexandrasamuel.com/dissertation/index.html.

Schneier, B. 2007. Gathering "Storm" Superworm Poses Grave Threat to PC Nets. *Counterpane* (04 October 2007), http://www.schneier.com/essay-184.html.

Schryen, G. 2007. *Anti-spam measures: analysis and design*: Springer-Verlag New York Inc.

Sendmail. 2007. Combatting Spam Best Practices. http://www.sendmail.com/sm/wp/spam_best_practices/.

Shadowserver. 2007a. Shadowserver Foundation - Botnet Information - What is a Botnet? , http://www.shadowserver.org/wiki/pmwiki.php/Information/Botnets.

———. 2007b. Shadowserver Foundation - Information - Honeypots. http://www.shadowserver.org/wiki/pmwiki.php/Information/Honeypots.

Shapiro, C., and H. Varian. 2000. *Information rules: A Strategic Guide to the Network Economy*: Harvard business school press Boston, MA.

Siegel, S., and N. Castellan. 1988. *Nonparametric statistics for the social sciences*: McGrawHill.

Sophos. 2007. Security threat report 2007. Sophos.

———. 2008. Security threat report 2008. Sophos.

Spamhaus. *The Spamhaus Project - The Definition Of Spam*. Available from http://www.spamhaus.org/definition.html.

———. 2009. A Snowshoe Winter: Our Discontent with CAN-SPAM. http://www.spamhaus.org/news.lasso?article=641.

TeamCymru. 2006a. Cybercrime: an epidemic. *ACM Queue* 4 (3):35.

———. 2006b. The Underground Economy: Priceless. *login* 31 (6).

TeleGeography. 2009. *Overview: GlobalComms Database: Voice Research Products: TeleGeography Research* 2009 [cited 25 October 2009]. Available from http://www.telegeography.com/products/global_comms/index.php.

Thurrott, P. 2004. Windows XP Service Pack 2 Beta Review. http://www.winsupersite.com/reviews/windowsxp_sp2_preview2.asp.

TrendLabs. 2009a. Internet Payment Site ClickandBuy Phished. *Malware Blog | Trend Micro*, http://blog.trendmicro.com/internet-payment-site-clickandbuy-phished.

———. 2009b. MSN Messenger Target For E-mail Phishing Attacks *Malware Blog | Trend Micro*, http://blog.trendmicro.com/see-who-blocked-you-on-msn-phishing-attacks.

TrendMicro. 2009. Trend Micro 2008 Annual Threat Roundup and 2009 Forecast.

UToronto. *Tracking GhostNet: Investigating a Cyber Espionage Network* 2009. Available from http://www.news.utoronto.ca/media-releases/international-affairs/information-warfare-monitor.html.

Van Eeten, M. J. G., and J. M. Bauer. 2008. Economics of Malware: Security Decisions, Incentives, and Externalities.

Velde, M. v. d., P. Jansen, and N. Anderson. 2007. *Guide to management research methods*.

Wash, R. 2007. Incentive design for home computer security.

Wikipedia. 2009. *DNSBL* 2009a [cited 18 October 2009]. Available from http://en.wikipedia.org/wiki/DNSBL.

———. 2009. *Domain name registrar*. Wikipedia, the fre encyclopedia 2009b [cited 09 October 2009]. Available from http://en.wikipedia.org/wiki/Domain_name_registrar.

———. 2009. *Economics of security* 2009c [cited 19 October 2009]. Available from http://en.wikipedia.org/wiki/Economics_of_security.

———. 2009. *Mail Abuse Prevention System*. Wikipedia 2009d [cited 25 October 2009]. Available from http://en.wikipedia.org/wiki/Mail_Abuse_Prevention_System.

———. 2010. *Comcast* 2010a [cited 20 January 2010]. Available from http://en.wikipedia.org/wiki/Comcast#Controversies.

———. 2010. *Network neutrality* 2010b [cited January 2010]. Available from http://en.wikipedia.org/wiki/Network_neutrality.

Zittrain, J. *ISPs helping with botnets* 2009. Available from http://www.interesting-people.org/archives/interesting-people/200910/msg00096.html.

APPENDIX A – SELECTED LAWS AND TREATIES

EU DIRECTIVE ON PRIVACY AND ELECTRONIC COMMUNICATIONS

DIRECTIVE 2002/58/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector

Article 13

Unsolicited communications

1. The use of automated calling systems without human intervention (automatic calling machines), facsimile machines (fax) or electronic mail for the purposes of direct marketing may only be allowed in respect of subscribers who have given their prior consent.

2. Notwithstanding paragraph 1, where a natural or legal person obtains from its customers their electronic contact details for electronic mail, in the context of the sale of a product or a service, in accordance with Directive 95/46/EC, the same natural or legal person may use these electronic contact details for direct marketing of its own similar products or services provided that customers clearly and distinctly are given the opportunity to object, free of charge and in an easy manner, to such use of electronic contact details when they are collected and on the occasion of each message in case the customer has not initially refused such use.

3. Member States shall take appropriate measures to ensure that, free of charge, unsolicited communications for purposes of direct marketing, in cases other than those referred to in paragraphs 1 and 2, are not allowed either without the consent of the subscribers concerned or in respect of subscribers who do not wish to receive these communications, the choice between these options to be determined by national legislation.

4. In any event, the practice of sending electronic mail for purposes of direct marketing disguising or concealing the identity of the sender on whose behalf the communication is made, or without a valid address to which the recipient may send a request that such communications cease, shall be prohibited.

5. Paragraphs 1 and 3 shall apply to subscribers who are natural persons. Member States shall also ensure, in the framework of Community law and applicable national legislation, that the legitimate interests of subscribers other than natural persons with regard to unsolicited communications are sufficiently protected.

THE LONDON ACTION PLAN

On International Spam Enforcement Cooperation

On October 11, 2004, government and public agencies from 27 countries responsible for enforcing laws concerning spam met in London to discuss international spam enforcement cooperation. At this meeting, a broad range of spam enforcement agencies, including data protection agencies, telecommunications agencies and consumer protection agencies, met to discuss international spam enforcement cooperation. Several private sector representatives also collaborated in parts of the meeting.

Global cooperation and public private partnerships are essential to spam enforcement, as recognized in various international fora. Building on recent efforts in organizations like the Organisation for Economic Cooperation and Development (OECD) and the OECD Spam Task Force, the International Telecommunications Union (ITU), the European Union (EU), the International Consumer Protection Enforcement Network (ICPEN), and the Asia-Pacific Economic Cooperation (APEC), the Participants issue this Action Plan. The purpose of this Action Plan is to promote international spam enforcement cooperation and address spam related problems, such as online fraud and deception, phishing, and dissemination of viruses. The Participants also open the Action Plan for participation by other interested government and public agencies, and by appropriate private sector representatives, as a way to expand the network of entities engaged in spam enforcement cooperation.

A. The participating government and public agencies (hereinafter "Agencies"), intend to use their best efforts, in their respective areas of competence, to develop better international spam enforcement cooperation, and intend to use their best efforts to:

1) Designate a point of contact within their agency for further enforcement communications under this Action Plan.

2) Encourage communication and coordination among the different Agencies that have spam enforcement authority within their country or region to achieve efficient and effective enforcement, and to work with other Agencies within the same country or region to designate a primary contact for coordinating enforcement cooperation under this Action Plan.

3) Take part in periodic conference calls, at least quarterly, with other appropriate participants to:

a) Discuss cases.
b) Discuss legislative and law enforcement developments.
c) Exchange effective investigative techniques and enforcement strategies.
d) Discuss obstacles to effective enforcement and ways to overcome these obstacles.
e) Discuss undertaking, as appropriate, joint consumer and business education projects addressing problems related to spam such as online fraud and deception, phishing, and dissemination of viruses. Such projects could include educational efforts addressing conditions facilitating the anonymous delivery of spam, such as the use of open relays, open proxies and zombie drones.
f) Participate as appropriate in joint training sessions with private sector representatives to identify new ways of cooperating and to discuss spam investigation techniques.

4) Encourage dialogue between Agencies and appropriate private sector representatives to promote ways in which the private sector can support Agencies in bringing spam cases and pursue their own initiatives to fight spam.

5) Prioritize cases based on harm to victims when requesting international assistance.

6) Complete the OECD Questionnaire on Cross border Enforcement of Anti Spam Laws, copies of which may be obtained from the OECD Secretariat.

7) Encourage and support the involvement of less developed countries in spam enforcement cooperation.

The participating Agencies intend to keep information shared in the context of this Action Plan confidential when requested to do so, to the extent consistent with their respective laws. Similarly, the participating Agencies retain the right to determine the information they share under this Action Plan.

B. The participating private sector representatives (whether as a group or through its members) intend to use their best efforts to develop public private partnerships against spam and to:

1) Designate a single spam enforcement contact within each organization, who would coordinate with spam enforcement agencies on requests for enforcement related assistance

2) Work with other private sector representatives to establish a resource list of individuals within particular sectors (e.g., Internet service providers, registrars, etc.) working on spam enforcement.

3) Participate as requested and appropriate in segments of the periodic conference calls described in paragraph A.3 above for the purpose of assisting law enforcement agencies in bringing spam cases. (Because some calls will be focused solely on law enforcement matters, private sector representatives will participate only in selected calls.) In these conference calls, the participating private sector representatives intend to use their best efforts to:

a. Report about:
    i) Cases involving spam or related matters.
    ii) New technology and trends in email and spam.
    iii) New ways of cooperating with Agencies.
    iv) Obstacles to cooperation with Agencies and within the private sector.
    v) General data on spam and on line fraud as an early warning mechanism for Agencies.
b. Assist as appropriate in training sessions on subjects such as the latest spam investigation techniques to help Agencies in investigating and bringing spam cases.

In order to prevent inappropriate access to information, a private sector representative may be excluded from participating in all or a portion of the periodic conference calls described above if a participating Agency objects.

4) Work cooperatively with Agencies to develop the most efficient and effective ways to frame requests for information. For this purpose, each participating private sector representative intends to use best efforts to compile written responses to the following questions:

a. What kind of information do you provide about potential spammers to domestic law enforcement agencies and under what circumstances?
b. What kind of information would you provide about potential spammers to foreign law enforcement agencies and under what circumstances?
c. How do you recommend that spam enforcement agencies submit requests for assistance to you?

C. In order to begin work pursuant to this Action Plan, the U.K. Office of Fair Trading and the U.S. Federal Trade Commission intend to use best efforts to:

1. Collect and disseminate information provided pursuant to this Action Plan, including points of contact, notifications from new Participants of their willingness to endorse this Action Plan, and responses to questionnaires, in cooperation with the OECD.

2. Set up the conference calls mentioned in paragraph A.3.

3. Provide a contact for further communications under this Action Plan.

The participating Agencies expect that this procedure may be modified at any time.

D. This Action Plan reflects the mutual interest of the Participants in the fight against illegal spam. It is not intended to create any new legally binding obligations by or amongst the Participants, and/or require continuing participation.

Participants to this Action Plan recognize that cooperation pursuant to this Action Plan is subject to their laws and their international obligations, and that nothing in this Action Plan requires the Participants to provide confidential or commercially sensitive information.

Participants in this Action Plan intend to use best efforts to share relevant findings of this group with the OECD Spam Task Force and other appropriate international groups.

This Action Plan is meant to be a simple, flexible document facilitating concrete steps to start working on international spam enforcement cooperation. It is expected that the collective work program under this Action Plan may be refined, and if necessary changed by the participants, as new issues arise.

Additional Agencies, and private sector representatives as defined below, may endorse and take part in this Action Plan as long as no Agency that has endorsed this Action Plan objects.

"Private sector representatives" invited to participate in this Action Plan include financial institutions, Internet service providers, telecommunications companies, information security software providers, mobile operators, courier services, commercial mail receiving agencies, industry membership organizations, consumer organizations, payment system providers, credit reporting agencies, domain name registrars and registries, and providers of alternative dispute resolution services.

For more information see: http://www.londonactionplan.com/?q=node/4

THE CONVENTION ON CYBERCRIME

The Convention on Cybercrime is the first international treaty seeking to address Computer crime and Internet crimes by harmonizing national laws, improving investigative techniques and increasing cooperation among nations.[1][2] It was drawn up by the Council of Europe in Strasbourg with the active participation of the Council of Europe's observer states Canada, Japan and USA.

The Convention and its Explanatory Report was adopted by the Committee of Ministers of the Council of Europe at its 109th Session on 8 November 2001. It was opened for signature in Budapest, on 23 November 2001 and it entered into force on 1 July 2004.[3] As of 2 September 2006, 15 states had signed, ratified and acceded to the convention, while a further 28 states had signed the convention but not ratified it.[4]

On 1 March 2006 the Additional Protocol to the Convention on Cybercrime came into force. Those States that have ratified the additional protocol are requited to criminalize the dissemination of racist and xenophobic material through computer systems, as well as of racist and xenophobic-motivated threats and insults.[5]

The Convention is the first international treaty on crimes committed via the Internet and other computer networks, dealing particularly with infringements of copyright, computer-related fraud, child pornography and violations of network security. It also contains a series of powers and procedures such as the search of computer networks and Lawful interception.

Objectives

Its main objective, set out in the preamble, is to pursue a common criminal policy aimed at the protection of society against cybercrime, especially by adopting appropriate legislation and fostering international co-operation.

The Convention aims principally at:

1. harmonising the domestic criminal substantive law elements of offences and connected provisions in the area of cyber-crime

2. providing for domestic criminal procedural law powers necessary for the investigation and prosecution of such offences as well as other offences committed by means of a computer system or evidence in relation to which is in electronic form

3. setting up a fast and effective regime of international co-operation.

The following offences are defined by the Convention: illegal access, illegal interception, data interference, system interference, misuse of devices, computer-related forgery, computer-related fraud, offences related to child pornography and offences related to copyright and neighbouring rights.

It also sets out such procedural law issues as expedited preservation of stored data, expedited preservation and partial disclosure of traffic data, production order, search and seizure of computer data, real-time collection of traffic data, and interception of content data. In addition, the Convention contains a provision on a specific type of transborder access to stored computer data which does not require mutual assistance (with consent or where publicly available) and provides for the setting up of a 24/7 network for ensuring speedy assistance among the Signatory Parties.

The Convention is the product of four years of work by European and international experts. It has been supplemented by an Additional Protocol making any publication of racist and xenophobic propaganda via

computer networks a criminal offence. Currently, cyber terrorism is also studied in the framework of the Convention.

Sources:
http://en.wikipedia.org/wiki/Convention_on_Cybercrime
http://conventions.coe.int/Treaty/EN/Treaties/Html/185.htm