Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development **12-Nov-2010**

_____

English - Or. English

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY

THE ROLE OF INTERNET SERVICE PROVIDERS IN BOTNET MITIGATION: AN EMPIRICAL ANALYSIS BASED ON SPAM DATA

STI WORKING PAPER 2010/5

By Michel van Eeten (Delft University of Technology), Johannes M. Bauer (Michigan State University), Hadi Asghari (Delft University of Technology), Shirin Tabatabaie (Delft University of Technology)

JT03292276

STI Working Paper Series

The Working Paper series of the OECD Directorate for Science, Technology and Industry is designed to make available to a wider readership selected studies prepared by staff in the Directorate or by outside consultants working on OECD projects. The papers included in the series cover a broad range of issues, of both a technical and policy-analytical nature, in the areas of work of the DSTI. The Working Papers are generally available only in their original language – English or French – with a summary in the other.

Comments on the papers are invited, and should be sent to the Directorate for Science, Technology and Industry, OECD, 2 rue André-Pascal, 75775 Paris Cedex 16, France.

The opinions expressed in these papers are the sole responsibility of the author(s) and do not necessarily reflect those of the OECD or of the governments of its member countries.

www.oecd.org/sti/working-papers

THE ROLE OF INTERNET SERVICE PROVIDERS IN BOTNET MITIGATION
AN EMPIRICAL ANALYSIS BASED ON SPAM DATA

ABSTRACT

Botnets – networks of machines infected with malicious software – are widely regarded as a critical security threat. Measures that directly address the end users who own the infected machines are useful, but have proven insufficient to reduce the overall problem. Recent studies have shifted attention to Internet Service Providers (ISPs), the providers of Internet access to end users, as possible control points for botnet activity.

In the report, we set out to empirically answer the following questions:

- First, to what extent are ISPs critical control points for botnet mitigation?

- Second, to what extent do they perform differently relative to each other, in terms of the number of infected machines in their networks?

- Third, and last, to what extent can we explain the differences in performance from the characteristics of the ISPs or the environment in which they are located?

We have gathered data on the location of infected machines over time by studying spam traffic. Around 80- 90 % of all spam is issued by infected machines. The origin of a spam message therefore very likely indicates the presence of an infected machine. Our raw data is a global dataset that comprises 109 billion spam messages from 170 million unique IP addresses, all of which were delivered to a 'spam trap' in the period 2005-2009.

Our findings lend direct and indirect support to the view that ISPs are important potential control points. The 200 ISPs that hold the lion's share of the access markets in a wider OECD area – by which we mean the 33 members, plus two "accession candidates" (Estonia and the Russian Federation) and the five "enhanced-engagement" countries (Brazil, China, India, Indonesia and South Africa) – harbor over 60 % of all infected machines worldwide registered by the spam trap. Furthermore, we discovered that infected machines display a highly concentrated pattern. The networks of just 50 ISPs account for around half of all infected machines worldwide. This is remarkable, in light of the tens of thousands of entities that can be attributed to the class of ISPs. The bulk of the infected machines are not located in the networks of obscure or rogue ISPs, but in those of established, well-known ISPs.

Not only do the legitimate ISPs harbor a large share of all infected machines, they also vary widely in their performance, which suggests that some have adopted more effective practices than others, even when operating under similar market and regulatory conditions. ISPs of similar size, operating in the same country, can differ by a factor of ten in the number of infected machines in their networks. While the strategies of the attackers are dynamic, the security performance of ISPs turns out to be quite stable. For the past four years (2006-2009), we looked at the composition of the group of 50 ISPs that had, in absolute terms, the most infected machines in their network. We found that 31 ISPs are in that top 50 in all 4 years. Geographically, they are distributed across 17 countries in the geographic area mentioned above.

For the first time, the patterns in infected machines are connected to other data, such as the size of the ISPs and the country in which they are located. Using bivariate and multivariate statistical approaches, the analysis finds that several variables are significant factors, such as: the size of the ISP matters, as do the

level of software piracy in a country and the activity of the regulator (as measured, for example, by participation in the London Action Plan).

Specific policy lessons have to be derived with caution and judgment. The findings reported in this study are based on past data and are only valid predictors of future events if the overall patterns continue to hold. The five years of observations seem to indicate, however, that despite new forms of malware and new attack strategies the overall emerging patterns are fairly robust.

LE ROLE DES FOURNISSEURS D'ACCES INTERNET POUR LA REDUCTION DES BOTNETS
UNE ANALYSE EMPIRIQUE FONDEE SUR LES DONNEES DU SPAM

RESUME

Les botnets, ces réseaux de machines infectées par du code logiciel malveillant, sont généralement considérés comme une menace critique pour la sécurité. Les mesures dirigées directement vers les utilisateurs finaux à qui appartiennent les machines infectées sont utiles, mais se sont montrées insuffisantes pour réduire le problème d'ensemble. Des études récentes ont réorienté l'attention vers les Fournisseurs d'Accès Internet (FAI), c'est-à-dire ceux qui connectent les utilisateurs finaux à Internet, en tant que possibles points de contrôle de l'activité des botnets.

Dans ce rapport, nous ébauchons une réponse empirique aux questions suivantes :

- Premièrement, dans quelle mesure les FAI sont-ils des points de contrôle critique pour la réduction des botnets ?

- Deuxièmement, dans quelle mesure sont-ils relativement plus performants les uns par rapport aux autres en termes de nombre de machines infectées dans leur réseau ?

- Troisièmement, dans quelle mesure peut-on expliquer les différences de performance par les caractéristiques des FAI ou l'environnement dans lequel ils sont situés ?

Nous avons collecté des données sur la localisation de machines infectées en étudiant le trafic lié au spam. Environ 80 % à 90 % de tout le spam provient de machines infectées. L'origine d'un message de spam indique donc par conséquent très probablement la présence d'une machine infectée. Nos données brutes constituent une base de données mondiale qui comprend 109 milliards de messages de spam en provenance de 170 million d'adresses IP uniques, l'ensemble desquels ont été envoyés à un « piège à spam » durant la période 2005-2009.

Nos conclusions soutiennent de façon directe et indirecte l'hypothèse selon laquelle les FAI sont d'importants points de contrôle potentiels. Les 200 FAI qui se partagent la part du lion du marché de l'accès Internet dans une zone OCDE élargie – c'est-à-dire les 33 pays membres plus deux pays candidats à l'adhésion (Estonie et Fédération de Russie) et les cinq pays de l'engagement renforcé (Afrique du Sud, Brésil, Chine, Inde, Indonésie) – comprennent plus de 60 % de toutes les machines infectées du monde. De plus, nous avons découvert que les machines infectées révèlent un schéma de forte concentration : le réseau de quelques 50 FAI regroupe environ la moitié de toutes les machines infectées dans le monde. Ceci est remarquable à la lumière des dizaines de milliers d'entités qui peuvent entrer dans la catégorie des FAI. Le gros des machines infectées n'est pas situé dans les réseaux de FAI obscurs ou voyous mais dans ceux des FAI établis et bien connus.

Non seulement les FAI légitimes accueillent une large part de toutes les machines infectées, mais leur performance varie également beaucoup, ce qui suggère que certains ont adopté des pratiques plus efficaces que d'autres, même opérant dans des conditions de marché et de régulation similaires. Le nombre de machines infectées dans le réseau de FAI de taille similaire opérant dans le même pays peut être jusqu'à 10 fois supérieur d'un fournisseur à l'autre. Alors que les stratégies des attaquants sont dynamiques, les performances de sécurité des FAI s'avèrent être assez stables. Sur les quatre dernières années (2006-2009), nous avons observé la composition du groupe des 50 FAI qui avaient, en valeur absolue, le plus grand

nombre de machine infectées dans leur réseau : 31 FAI sont dans ce top 50 durant les 4 années. Géographiquement, ils sont répartis dans 17 pays de la zone géographique mentionnée ci-dessus.

Pour la première fois, les modèles d'infections de machines sont reliés à d'autres données, telles que la taille des FAI et le pays dans lequel ils sont situés. En utilisant des approches statistiques bivariées et multivariées, l'analyse montre que plusieurs variables sont des facteurs significatifs tels que : la taille du FAI, le niveau de piratage dans le pays ainsi que l'activité de régulation (telle que mesurée par la participation dans le Plan d'Action de Londres).

Il est nécessaire de procéder avec prudence et discernement pour tirer des enseignements spécifiques pour les politiques publiques. Les conclusions de cette étude sont fondées sur des données passées et ne permettent de prédire le futur que si les schémas d'ensemble persistent. Les cinq années d'observations semblent cependant indiquer que malgré les nouvelles formes de logiciels malveillants et les nouvelles stratégies d'attaque, les schémas d'ensemble qui émergent sont assez robustes.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Botnets – networks of machines infected with malicious software – are widely regarded as a critical security threat. While originating in criminal behaviour, the magnitude and impact of botnets is also influenced by the decisions and behaviour of legitimate market players in the Internet economy, such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users.

Focus on Internet Service Providers

Measures that directly address the end users who own the infected machines are useful, but have proven insufficient to reduce the overall problem. Recent studies have shifted attention to Internet Service Providers (ISPs), the providers of Internet access to end users, as possible control points for botnet activity.

There are three, often overlooked, assumptions behind the focus on ISPs as key intermediaries for botnet mitigation. First, it assumes that ISPs are a critical control point for infected machines. This may seem obvious, but it has never been empirically established what portion of infected machines on the Internet is actually located within the networks of ISPs – as opposed to, say, hosting providers, application service providers, webmail providers, university networks and corporate networks.

The second assumption behind the focus on ISPs, is that the ISPs who would carry the burden of increased mitigation efforts are also the ones who control the bulk of the problem. We are most familiar with the legitimate ISPs – well-known brands that together possess the bulk of the market share. These organisations are identifiable, reachable and stable enough to be brought into some form of collaboration or under a regulatory regime. Treating ISPs as control points implicitly assumes that the problem exists for the most part within the networks of these providers; not in the margins of the market, which is teeming with large numbers of small ISPs, among which are the so-called 'rogue ISPs' that facilitate or at least condone criminal activity. These small ISPs are often shortlived and difficult to survey, let alone reach through collaborative efforts or public regulation. They also typically evade, intentionally or not, the collaborative processes through which collective action is brought about.

The third assumption is that ISPs have discretion over their mitigation efforts. In other words, the incentives under which they operate allow them to increase their efforts. It is not clear to what extent ISPs are constrained by their market and institutional environment in freely determining their own policies in this regard. If, for example, the market for Internet access is characterised by fierce price competition, ISPs would be strongly discouraged to invest more in botnet mitigation than their competitors – *i.e.*, they would be disincentivised to contact and quarantine more infected customers than their competitors. These conditions would force ISPs to perform at similar levels in terms of security. If ISPs have only very limited discretion, then any attempt to increase the performance would have to change the conditions under which ISPs operate – at least for enough of them to reach critical mass. Voluntary initiatives for botnet mitigation measures are less likely to succeed.

Do the three assumptions outlined above hold in reality? If not, then there is reason to doubt the effectiveness of the current initiatives to stimulate ISPs to step up their botnet mitigation efforts. With some effort, all three assumptions can be tested empirically. The first two can be tested jointly by

identifying what portion of the overal population of infected machines is located within the networks of the established ISPs that hold the bulk of the market share. The third assumption can be tested by comparing the levels of infections of ISPs that operate under similar conditions. If these levels are substantially different, it suggests that ISPs have discretion and that their security policies make a difference.

Research approach

In the report, we set out to empirically answer the following questions:

- First, to what extent are ISPs critical control points for botnet mitigation?

- Second, to what extent do they perform differently relative to each other, in terms of the number of infected machines in their networks?

- Third, and last, to what extent can we explain the differences in performance from the characteristics of the ISPs or the environment in which they are located?

We have gathered data on the location of infected machines over time by studying spam traffic. Around 80 to 90 % of all spam is issued by infected machines. The origin of a spam message therefore very likely indicates the presence of an infected machine.

Our raw data is a global dataset that comprises 109 billion spam messages from 170 million unique IP addresses, all of which were delivered to a 'spam trap' in the period 2005-2009. We have extensively tested the representativeness of this data and found that it is consistent with the data published by security service providers. Next, we developed a variety of strategies to compensate for known limitations when using IP addresses to count machines connected to the Internet. Using the spam data, we have analyzed in detail the geographic patterns, time trends, and differences at the level of countries and ISPs.

Key findings

Our findings lend direct and indirect support to the view that ISPs are important potential control points. The 200 ISPs that hold the lion's share of the access markets in a wider OECD area – by which we mean the 33 members, plus two "accession candidates" (Estonia and the Russian Federation) and the five "enhanced-engagement" countries (Brazil, China, India, Indonesia and South Africa) [1] – harbor over 60 % of all infected machines worldwide. Furthermore, we discovered that infected machines display a highly concentrated pattern. The networks of just 50 ISPs account for around half of all infected machines worldwide. This is remarkable, in light of the tens of thousands of entities that can be attributed to the class of ISPs. The bulk of the infected machines are not located in the networks of obscure or rogue ISPs, but in those of established, well-known ISPs.

Not only do the legitimate ISPs harbor a large share of all infected machines, they also vary widely in their performance, which suggests that some have adopted more effective practices than others, even when operating under similar market and regulatory conditions. ISPs of similar size, operating in the same country, can differ by a factor of ten in the number of infected machines in their networks. Comparing ISPs across countries, we observe differences of two orders of magnitude, *i.e.* a hundred times more infections in networks of similar size. These differences in security performance imply that ISPs have leeway to increase their efforts, that security performance is not dictated by external conditions.

Some ISPs already contact and, when needed, quarantine infected customers. While those efforts likely reduce their infection rates, they are typically limited when compared to the overall scale of the

problem. There is no publicly available data about how many customers are contacted by ISPs, but when we compare anecdotal evidence to our estimates of the number of infected machines in the networks of ISPs, it suggests that even good ISPs contact only a fraction of the affected customers.

While the strategies of the attackers are dynamic, the security performance of ISPs turns out to be quite stable. We looked in more detail at ISPs in the group of poor performers. For the past four years (2006-2009), we looked at the composition of the group of 50 ISPs that had, in absolute terms, the most infected machines in their network. We found that 31 ISPs are in that top 50 in all 4 years. These ISPs range from the large to the relatively small – from 197 000 to 53.5 million subscribers. Geographically, they are distributed across 17 countries in the geographic area described above. When looking at relative infection rates – *i.e.* the number of infections per subscriber – similar patterns emerged.

We derived metrics, such as the average number of infected machines per subscriber, that were then used as dependent variables to be explained by predictors related to the institutional framework, operational characteristics of ISPs, and other control variables. With regard to the control variables, we found that characteristics of the user base matter. In countries where consumers are more likely to use pirated software, we find higher botnet activity. The level of education, as a proxy for technical competence, is associated with lower levels of botnet activity. Higher average connection speeds are, however, not associated with higher levels of botnet infections, as is often presumed. In fact, we find the reverse: that high connection speeds are associated with lower botnet activity. A number of other control variables, such as the income level of a country, did not turn out to be significant factors in explaining the average number of infected sources per subscriber.

We found evidence to support the idea that broad governmental efforts to improve cybersecurity are associated with lower levels of botnet activity. Membership of the London Action Plan (LAP) was used as a proxy for how active regulatory authorities have been in the area of cybersecurity. Across the variety of statistical models we developed, LAP membership was consistently associated with lower infection rates among ISPs in those countries. This evidence must not be mistaken as a causal relation in the sense that joining LAP alone will reduce botnet activity. Rather, it suggests that organisation membership contributes to other government and non-government measures that have mitigating effects on botnets. Qualitative evidence from selected countries supports this finding, as governements known for their active engagement of ISPs in the area of security – most notably, Japan and Finland – display better performance. Our data also confirm that the infection rates of the ISPs in those countries are among the lowest worldwide.

We also tested several of the factors that have been considered as important in explaining ISP security performance. Average revenue per customer did not make a difference, which leads us to conclude that the competitive intensity of ISPs' market environment does not have a direct influence on security performance. We also tested the claim that large ISPs perform worse than smaller ones. Some experts have argued that large ISPs are less subject to peer pressure. Our data suggests that this is incorrect. In fact, large ISPs perform slightly better than average (measured by the number of infected sources and spam volume per subscriber). The market share of an ISP in its home country was not associated with worse performance either.

One educated guess as to why large ISPs actually do slightly better, is that their size forces them to introduce automation in incident response and abuse management. A similar mechanism may explain why we found that cable providers did better than DSL providers, especially among smaller ISPs. Management of cable networks often relies on automated systems and these technologies might reduce the cost of dealing with infected machines. Given the ongoing advances in technology, including botnet mitigation solutions, the difference between cable and DSL may disappear in the immediate future. Our findings do imply, however, that automation is likely to be a critical part of scaling up ISP efforts.

Policy implications

From a policy perspective, the finding that a relatively small number of ISPs is associated with a large share of total spam activity is relevant. Although these ISPs are not themselves the origin of botnet infections, they play an important role in the chain from cybercriminals to the targets of botnet attacks. The study uncovered a specific pattern that suggests that the chances of devising meaningful forms of private and public sector measures might be higher than commonly thought. The highly concentrated pattern we uncovered suggests that the number of actors needed to create an impact on botnets is smaller than expected. It would be extremely difficult to bring about collective action among many thousands of ISPs located in over a hundred countries, even if ISPs were to be a more effective control point than the billion to billion-and-a-half end users. Furthermore, the most critical actors are larger, well-established corporations. It may be easier to design public policy measures and implement them for this group of ISPs, whether such measures are government interventions or forms of public-private sector co-operation. Such measures would be much more difficult if large numbers of small ISPs that are often shortlived and difficult to survey were involved. Many of these small ISPs are difficult to reach with collaborative or regulatory efforts. Even if they were interested in co-operation, the transaction cost of bringing large numbers of players into the fold may be very high.

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

While ISPs appear to have considerable discretion, their incentives are also shaped by external conditions. In retail ISP markets competition is primarily driven by price and in many countries price competition is fierce. Even if price does not seem to have a significant influence on security performance, from an ISP's point of view, policy measures that affect costs (and all do directly and indirectly) are unfunded mandates and may be difficult to realise given this competitive environment. Thus, it may be necessary to think about innovative funding schemes for such programmes. Moreover, even if consumers cared about security, there are no adequate market signals that could reliably guide them towards better performing ISPs. Establishing a trusted rating system might be a tool to assist consumers in this regard. Most industry insiders lack such signals as well, except for the anecdotal evidence and speculative claims about the performance of this or that ISP that are bandied among the members of the security community. Current efforts to bring about collective action – through industry self-regulation, co-regulation, or government intervention – might initially achieve progress by focusing on the set of ISPs that together have the lion's share of the market.

Specific policy lessons have to be derived with caution and judgment. Malware and botnets are dynamic phenomena. History tells us that every fortification of information security will trigger adaptations in attack strategies. Likewise, any reduction of the intensity of attacks may tempt users to reduce security investment. Both effects imply that the emergent level of security at the sytem level may respond less to policy measures than hoped. Our data point to considerable inertia in the system, which could be seen as one outcome of these effects. The findings reported in this study are based on past data and are only valid predictors of future events if the overall patterns continue to hold. The five years of

observations seem to indicate, however, that despite new forms of malware and new attack strategies the overall emerging patterns are fairly robust. Lastly, while the study is based on detailed data, lots of information that would be required to formulate coherent and effective policies is not systematically collected or not in the public domain. This is particularly true for information on damages from breaches of information security and for data on ISP-level security measures that would help assess which firm-level strategies are effective. The findings of the study need to be interpreted with these caveats in mind.

THE ROLE OF INTERNET SERVICE PROVIDERS IN BOTNET MITIGATION
AN EMPIRICAL ANALYSIS BASED ON SPAM DATA[2]

INTRODUCTION

Malicious software has been used to infect tens of millions of computers worldwide. Many of these machines are recruited into so-called botnets; networks of thousands or even millions of computers that can be used for criminal purposes. This report addresses the issue of how the threat posed by botnets can be best mitigated. The introduction discusses why the focus in mitigation has shifted from end users, who own most of the infected machines, to Internet Service Providers (ISPs). It ends with articulating three research questions that set out to test the empirical merits of this focus on ISPs. The second part describes in detail the methodology that was developed to answer these questions. The third and final part discuss the empirical findings as well as their policy implications.

Botnets: A security threat to the Internet economy

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

If recent estimates are correct, around 5 % of all machines connected to the Internet may be infected with malware (BBC News 2007; House of Lords 2007; Moore *et al.* 2009; Clayton 2010). The fact that the owners of these machines often do not know their machines are compromised is part of the problem. Malware may be distributed and used in many ways, including e-mail messages, USB devices, infected websites, malicious advertising, and browser vulnerabilities (Jakobsson and Zulfikar, 2008).

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

Individuals, organisations, and nation states are targets. Attack vectors directed toward individuals and organisations include spam (most originating from botnets), variations of socially engineered fraud such as phishing and whaling, identity theft, attacks on websites, corporate espionage, "click fraud", "malvertising", and corporate espionage. The DDoS attacks on Estonia in 2007 and the spread of the

cryptic Conficker worm that, early in 2009, paralysed parts of the British and French military as well as government and health institutions in other countries, are recent examples of attacks on nations and their civil and military infrastructures (Clover 2009; Soper, 2009).

Security incentives of legitimate market players

While originating in criminal behaviour, the magnitude and impact of botnets is also influenced by the decisions and behaviour of legitimate market players in the Internet economy, such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users.

As security comes at a cost, tolerating some level of insecurity is economically rational from an individual and social point of view. Although it is mostly provided by private players, information security has strong public good characteristics. Therefore, from a societal perspective, a crucial question is whether the costs and benefits taken into account by market players reflect the social costs and benefits. If that is the case, decentralised individual decisions also result in an overall desirable outcome (*e.g.* a tolerable level of cybercrime, a desirable level of security). However, if some of the costs are borne by other stakeholders or some of the benefits accrue to other players (*i.e.* they are "externalised"), individual security decisions do not properly reflect social benefits and costs. This is the more likely scenario in highly interdependent information and communication systems such as the Internet.

Whereas the security decisions of a market player regarding malware will be rational for that player, given the costs and benefits it perceives, the resulting course of action inadvertently or deliberately imposes costs on other market players and on society at large. Decentralised individual decisions will therefore not result in a socially optimal level of security. The presence of externalities can result in Internet-based services that are less secure than is socially desirable.

The interdependence between stakeholders has, to a certain degree, contributed to a blame game between individual players accusing each other for insufficient security efforts. Whereas Microsoft, due to its pervasive presence, is a frequent target (*e.g.*, Perrow 2007), the phenomenon is broader: security-conscious ISPs blame rogue ISPs, ISPs blame software vendors, software vendors blame end users, countries with higher security standards blame those with lacking law enforcement, and so forth. Although these claims are sometimes correct, they are overly broad and simplistic. Empirical field work into the incentives of these players, makes it clear that the relationships are much more complicated than these accusations suggest (cf. Anderson *et al*. 2008; Van Eeten and Bauer 2008). All players operate under mixed incentives. Some security-enhancing behaviours are rewarded, others discouraged. This is not unique for any specific type of market player.

Shifting focus from end users to Internet intermediaries

Botnets are typically associated with the infected personal computers (PCs) of end users. But malware can be used to infect all kinds of Internet-connected machines and recruit them into botnets. In addition to PCs, botnets have been found to consist of web servers, routers, mobile devices and recently even SCADA systems, the "supervisory control and data acquisition" machines used to control energy systems and other infrastructural equipment.

Notwithstanding this diversity, a large portion of the machines in botnets are assumed to belong to end users, in particular home users and small and medium-size enterprise (SME) users. Compared to larger corporate users, these groups often do not achieve adequate levels of protection (Van Eeten and Bauer, 2008).

Security measures that address end users directly – including awareness raising and information campaigns – are useful, but they have proven to be insufficient to reduce the overall problem. Not because end users are incorrigible. Some surveys suggest that they do, in fact, increasingly adopt more secure practices, such as using anti-virus protection, a firewall, and automatic security updates for their software (*e.g.* Fox 2007). The attackers, however, also adapt and innovate their strategies. The net result is an inadequate defense against malware infections: while the capabilities and practices of end users are improving, they lag behind the increasingly sophisticated threats of attackers.

Many experts have pointed out the flaws associated with a focus on end users. In 2007, a review by the U.K. House of Lords (2007, p. 80) summarised the now dominant critique of this approach: "The current emphasis of Government and policy-makers upon end-user responsibility for security bears little relation either to the capabilities of many individuals or to the changing nature of the technology and the risk."

In light of these shortcomings, the focus has recently shifted from end users to other market players. This is not to say that end users are now absolved of responsibility. The shift merely signals that the efforts of other players are also necessary for an adequate response against botnets. Special attention is being paid to Internet intermediairies – market players who could function as natural "control points" for end user activities.[3] High on the list of intermediairies that are relevant to the fight against botnets, are the ISPs. Since ISPs – in the sense of access providers – are the gateway between their customers and the wider Internet, they are in a unique position to detect and mitigate the malicious activity of their customers' machines. An increasing number of proposals and initiatives are launched aimed at mobilising ISPs in the fight against botnets.

Botnet mitigation by Internet Service Providers

The idea that ISPs are in a position to mitigate the impact of botnets is broadly supported. Of course, the fact that ISPs *can* mitigate this threat, does not mean that they *should* mitigate it. They are not the origin of the externality and they have to bear substantial direct and indirect costs if they do internalise the externalities of their customers.

Nevertheless, in a variety of countries, ISPs are now explicitly assuming at least partial responsibility for botnet mitigation. Industry collaborative efforts like the Internet Engineering Taskforce (IETF) and the Messaging Anti-Abuse Working Group (MAAWG) have prepared sets of best practices for the remediation of bots in ISP networks. Under pressure from the government, Australia's largest ISPs have agreed on a voluntary code of conduct that includes contacting infected customers and filtering their connection.

Within the OECD, other countries have indicated they are pursuing similar lines of action. A related initiative in Germany is the establishment of a government-funded call centre to which ISPs can direct customers in need of support to disinfect their machines. In Japan, industry collaborates with the government in the Cyber Clean Center, which offers free botnet removal software directly to the infected users of participating ISPs. The main ISPs in the Netherlands – with an aggegate share of over 95 % of the retail market – have entered into a covenant that expresses their commitment to mitigate botnet activity in their own networks.

In these and other countries, many ISPs claim that their organisations already have practices in place where they contact and in some cases quarantine customers whose machines are infected with malware. While this may be true, there is currently no data available that indicates the scale on which these practices are being carried out.

Scale is critical, however. There are indications that ISPs only deal with a fraction of the infected machines in their networks. For example, in an earlier study we found that a large ISP with over 4 million customers contacted around 1 000 customers per month (Van Eeten and Bauer 2008). Typical estimates of security researchers put the number of infected machines at around 5 % of all connected machines at any point in time (Moore *et al*. 2009). This would translate into about 200 000 infected machines for this specific ISP. Even if we reduce the estimated infection rate to 1 %, that still implies 40 000 infected machines. This stands in stark contrast to the 1 000 customers that the ISP claimed to be contacting – even when we optimistically assume that all contacted customers are either willing and able to clean up their infected machine or are being quarantined.

Previous research would indeed predict such a discrepancy between ISP efforts and the actual number of infected machines. Some reports went as far as arguing that ISPs have "no incentive" to disconnect infected machines from their networks (House of Lords, 2007). Other studies claimed that small ISPs may have some "weak positive incentives", but that large ISPs "enjoy a certain impunity" and only "face limited economic incentive to clean up their act" (Anderson *et al.* 2008). Our own work is slightly more positive about the incentives of ISPs, but also signals significant areas where these are lacking or too weak to trigger security improvements (*e.g.*, Van Eeten and Bauer, 2008).

While there are differences among the findings of these studies, they do converge on the hypothesis that the incentives under which ISPs operate discourage them to scale up mitigation efforts to match the size of the problem. Currently, there is no authoritative economic research available to either refute or confirm this claim. A key factor is the cost of customer support. When an ISP contacts or quarantines infected customers, it will trigger incoming customer calls. The ISP incurs a certain cost to handle each call. Some ISPs have reported this cost to be around EUR 8 per incoming call, other estimates are substantially higher (;Van Eeten and Bauer, 2008, Clayton 2010). There are indications that the cost of support can quickly outweigh the profit margin for a subscription. Clayton recently estimated that two customer calls in a year may be enough to consume the profits on that customer (Clayton, 2010).

In light of these incentives, a few rather controversial proposals have emerged that try to move beyond the current voluntary efforts, asking governments to force ISPs to assume more responsibility. For example, Anderson *et al.* (2008) propose that liability for infected machines should be assigned to the ISPs, rather than to the consumers who own the machines. The authors also suggest to impose statutory damages on ISPs that do not respond promptly to requests for the removal of compromised machines. Another proposal has been to subsidise the clean-up of infected machines, in what the author calls a public health approach to cybersecurity (Clayton 2010). This also aims to overcome the incentive problem of ISPs, in this case by publicly funding the support cost of customers. The possible effects of these proposals are not yet fully understood. They will have to be investigated in more detail before these ideas can be transformed into realistic policy options.

To reiterate: We are not claiming that ISPs should contact all the owners of infected machines. That is a matter for policy development to consider, taking into account the costs and benefits of mitigation for ISPs, their customers, as well as society at large. We are simply stating that there is an urgent need to collect data, beyond the generic claims of ISPs that they are contacting customers and quarantining infected machines. Ideally, this data should inform us about the size of the problem, the extent to which ISPs can mitigate and are actually mitigating botnet activity, and how ISPs perform relative to each other. This reports sets out to address these knowledge gaps. Before we turn to the empirical research, however, it is important to better understand the assumptions underlying the shift in focus towards ISPs. That will help us frame the relevant research questions.

Assumptions underlying the focus on ISPs

There are three core assumptions behind the focus on ISPs as key intermediaries for botnet mitigation. First, it assumes that ISPs are a critical control point for infected machines. To some extent, this is obvious. The ISP's customers can only send and receive traffic via the ISP, which creates a natural bottleneck to mitigate malicious activity of the customers' machines. However, it has never been empirically established what portion of infected machines on the Internet is actually located within the networks of ISPs. What about the machines in use by, for example, hosting providers, application service providers, webmail providers, university networks and corporate networks? If ISPs can only control a minor portion of the infected machines, it undermines the argument to focus on them, more than on other players, as the key intermediaries in the fight against botnets.

The second assumption behind the focus on ISPs, is that the ISPs who would carry the burden of increased mitigation efforts are also those who control the bulk of the problem. This assumption is often overlooked in the debate. We are most familiar with the legitimate ISPs – well-known brands that together possess the bulk of the market share. These organisations are identifiable, reachable and stable enough to be brought into some form of collaboration or under a regulatory regime. Treating ISPs as control points implicitly assumes that the problem exists for the most part within the networks of these providers; not in the margins of the market, which is teeming with large numbers of small ISPs that are often shortlived and difficult to survey, let alone reach through collaborative efforts or public regulation. They also typically evade, intentionally or not, the collaborative processes through which collective action is brought about.

As security incidents have illustrated, the margin of the market is also where we find a class of so-called "grey" and "rogue" ISPs. If this class has a disproportionate share of the infected machines in itsnetworks, then incentivising ISPs to increase botnet mitigation through voluntarily or other types of policy measures may not work. The burden of such measures will likely not fall onto these grey and rogue ISPs, as they will evade these measures. Rather, it will affect mostly the legitimate, more established providers. In the worst case scenario, such initiatives would raise the operating cost of the established providers – and indirectly the price of Internet access – without getting increased security in return. Depending on the price elasticity of demand for Internet access, it may even lead to a shift of users to these fringe competitors and even worsen security. Thus, whether working with ISPs is a feasible policy option will critically depend on the share of infected machines on their networks.

The third assumption is that ISPs have discretion over their mitigation efforts. In other words, the incentives under which they operate allow them to increase their efforts. It is not clear to what extent ISPs are constrained by their market and institutional environment in freely determining their own policies in this regard. If, for example, the market for Internet access is charcterised by fierce price competition, ISPs would be strongly discouraged to invest more in botnet mitigation than their competitors – *i.e.* they would be disincentivised to contact and quarantine more infected customers than their competitors. These conditions would force ISPs to perform at similar levels in terms of security. Next to market incentives, institutional conditions could also reduce the discretion of ISPs. In some countries, ISPs reported that stringent data protection legislation severly limited the extent to which they could monitor their own networks and identify affected customers (Van Eeten and Bauer, 2008). If ISPs have only very limited discretion, then any attempt to increase the performance would have to change the conditions under which the ISPs operate. Voluntary initiatives for botnet mitigation initiatives are less likely to succeed.

Research questions

Do the three assumptions outlined above hold in reality? If not, then there is reason to doubt the effectiveness of the current initiatives to stimulate ISPs to step up their botnet mitigation efforts. With some effort, all three assumptions can be tested empirically. The first two can be tested jointly by

identifying what portion of the overal population of infected machines is located within the networks of the established ISPs that hold the bulk of the market share. The third assumption can be tested by comparing the levels of infections of ISPs that operate under similar conditions. If these levels are substantially different, it suggests that ISPs have discretion and that their security policies make a difference.

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

Before turning to these questions, subsequent sections of this report first outline the research approach as well as its limitations. At the heart of the research is data from a spam trap that has logged around 170 million unique IP addresses of machines that connected to it in the period 2005-2009. The raw data was parsed to associate IP addresses with ISPs and countries. We also discuss the limitations of the particular approach to mine the raw data and strategies to compensate for them. We then examine the intermediairy position of ISPs. Surprisingly, in our dataset, just 50 ISPs account for more than half of all infected machines worldwide. We also explore the differences among ISPs in the extent in which their networks harbor infected machines. These differences turn out to be substantial. To explain the differences, we employ bivariate and multivariate statistical approaches. Among others, using ISPs as the unit of analysis, we investigate empirically the effects of country-level policy measures on the number of infected machines sending spam. We conclude with a discussion of the implications of our findings for current efforts to mobilise ISPs in botnet mitigation.

METHODOLOGY

Identifying infected machines from spam data

There is currently no authoritative data source to identify the overal population of infected machines around the world. Commercial security providers typically use proprietary data and shield their measurement methods from public scrutiny. This makes it all but impossible to correctly interpret the figures they report and assess their validity.

The publicly accessible research in this area relies on two types of data sources:

i) *Data collected external to botnets*. This data identifies infected machines by their telltale behaviour, such as sending spam or participating in distributed denial of service attacks.

ii) *Data collected internal to botnets*. Here, infected machines are identified by intercepting communications within the botnet itself, for example by infiltrating the command and control infrastructure through which the infected machines get their instructions.

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

Neither type of data sources sees all infected machines, they only see certain subsets, depending on the specific data source. In general, one could summarise the difference between the first and the second source as a tradeoff between represenativeness versus accuracy. The first type captures a more representative slice of the problem, but will also include false positives. The second type accurately identifies infected machines, but only for a specific botnet, which implies that it cannot paint a representative picture.

This study draws upon a data source of the first type, namely spam traffic. The bulk of all spam messages are sent through botnets. Spammers use the thousands or millions of infected machines in a botnet to send out spam. Of the total volume of spam messages that are being sent out everyday, the overwhelming majority is sent through an infected machine. A variety of studies published during the period under study (2005-2009) found consistently that around 80–90 % of the total amount of spam comes from botnets (see Box 1). The IP address of the machine that delivered the spam message is therefore very likely to indicate the presence of an infected machine. Previous studies have also employed the origins of spam messages as proxy data to identify infected machines (*e.g.*, Zhuang *et al*. 2008).

Our data is drawn from a spam trap – an Internet domain set up specifically to capture spam, whose e-mail addresses have never been published or used to send or receive legitimate e-mail traffic. There is no legitimate way to deliver e-mail to the domain. All the e-mail it receives is indeed spam – as confirmed by logging the content of the messages. In the period of 2005-2009, the trap has received 109 billion spam messages from about 170 million unique IP adresses worldwide.

Box 1. Botnets are responsible for the bulk of all spam

It has been well established that for the period under study (2005-2009), the overwhelming majority of spam messages were issued through infected machines. Symantec's MessageLabs reported that in 2009, 83.4 % of all spam was issued by botnets.[4] Their estimates for the preceding years were: 90 % in 2008,[5] 80 % in 2006,[6] and "more than 80 percent" in 2005.[7] We could not find an estimate from MessageLabs for 2007. Cisco's Ironport, however, reported that in 2007, bots were responsible for 95 % of all spam[8] – an estimate that, according to the staff report, was supported by the nearly 50 panelists at the 2007 Spam Summit of the U.S. Federal Trade Commission.[9] As early as 2004, Sandvine reported that "up to 80 %" of all spam came from infected machines.[10]

A clear illustration of the large overlap between botnets and spam distribution came late 2008, when the shutdown of McColo, a U.S.-based Web host, led to an immediate drop of worldwide spam levels by half or even three-quarters, compared to previous levels.[11] McColo hosted the command and control servers for several botnets used for spam, such as a massive botnet called Srbizi. While most estimates are in the range of 80-90 %, we should also note that one firm, Sophos, has reported lower numbers. For 2005, it estimated that about 50 % of all spam came from botnets.[12] In its annual report for 2009, it simply stated that "the majority of spam is sent via botnets of hijacked systems in the homes and offices of innocent users".[13] However, it also noted that the shutdown of McColo, and thus of the Srbizi and other botnets, cut the global spam volume by three-quarters, which implies that its more recent estimates are closer to those of other firms.

Limitations of using spam data

Spam-sending machines are a powerful proxy for infected machines, but not a perfect one. As with any data source, spam data has certain limitations. Not all spam comes from infected machines and not all infected machines send spam.

The first issue points to the risk of false positives. Some machines that send spam are not infected machines. They could be hijacked webmail accounts or accounts with legitimate providers that have been bought by spammers for so-called "snowshoe spamming" campaigns.[14] For several reasons, we think our data is not significantly affected by this limitation. First, as was mentioned above, studies estimate that 80-90 % of all incoming spam originates from a botnet. We have reason to believe that for the spam received by our trap, this ratio is even higher. The trap is located at a small and relatively old generic top-level domain. It's e-mail addresses are therefore also relatively old. When spammers use different means of distribution than botnets, such as "snowshoe spamming", they tend to use targeted and fresher address lists, because these means of distribution are generally more costly than the use of botnets. In other words, this non-botnet spam would not be captured by our trap and not lead to false positives.

A more important mitigating measure against false positives is that we split all spam sources in two categories, depending on whether the network in which the source is located belongs to an ISP or not. Our analysis focuses mainly on the sources within ISP networks, which eliminates a lot of the potential false positives, namely spam from sources such as webmail providers, hosting providers and university networks. In sum, our approach implies that the impact of false positives is very limited.

The second issue – not all infected machines send spam – points to the risk of false negatives, of undercounting infected machines. Our data undoubtedly suffers from undercounting, as do all existing data sources on infected machines. That being said, sources external to botnets, such as spam traps, are less affected by this limitation than sources internal to botnets, such as data captured through infiltration of a specific botnet. While the latter can only see bots within a single botnet, the external sources can identify infected machines across a wide range of botnets – in our case, all botnets that send out spam, which is arguably the most common usage of botnets. In that sense, these sources can be considered the most representative of the overall population. Another factor that causes undercounting is that the spam logged by the trap is, of course, a sample of the global spam traffic. There are many infected machines that send spam, but only those messages that include the addresses of the spam trap among its intended destinations will show up in our data. For our research questions, this need not be a problem, since we are interested in relative metrics, *i.e.*, the infection rate of an ISP compared to other ISPs and to non-ISPs. A sample also allows this type of analysis, as long as it is representative for the whole population of spam messages. We have gone to great lengths to test the representativeness by comparing our data to the publicly available spam reports of commercial security providers. These tests confirm that the data is indeed representative (see Appendix 1).

Identifying the location of infected machines

Each machine that connects to the spam trap to deliver a spam message can be identified by its IP address. For each unique IP address that was logged by the spam trap, we looked up the Autonomous System Number (ASN) and the country where it was located. The ASN is relevant, because it allows us to identify what entity connects the IP address to the wider Internet – and whether that entity is an ISP or not. We looked up the country of an IP address by using so-called geo-IP data, which associates IP addresses with geographical locations – in this case, we used the MaxMind geoIP database.

As both ASN and geoIP information change over time, we used historical records to establish the orgin for the specific moment in time at which the message was received. We also recorded the number of spam messages sent from each source.[15] This effort resulted in two time series of variables: unique IP addresses and spam volume, both per ASN and per country. The number of unique IP addresses is more directly related to the number of infected machines. The latter variable is useful to balance some of the shortcomings of the former – a point to which we return later.

We then set out to identify which of the ASNs from which the trap received spam, belonged to ISPs. To the best of our knowledge, there is no existing database that maps ASNs onto ISPs. This is not surprising. Estimates of the number of ISPs vary from around 4 000 – based on the number of ASNs that provide transit services – to as many as 100 000 companies that self-identify as ISPs – many of whom are virtual ISPs or resellers of other ISPs' capacity.[16]

So we adopted a variety of strategies to connect ASNs to ISPs. First, we used historical market data on ISPs – wireline, wireless and broadband – from TeleGeography's GlobalComms database. We extracted the data on all ISPs in the database listed as operating in a set of 40 countries, namely all 34 members of the Organisation for Economic Co-operation and Development (OECD), plus one "accession candidate" and five so-called "enhanced-engagement" countries.

Box 2. How to detect infected machines via a spam trap

The bulk of all spam comes from botnets (see Box 1). By identifying where spam comes from, we can track the location of infected machines. This location can be identified by its IP address.

We captured a sample of the global spam traffic with a spam trap. This is a mail server that only has one purpose: to receive spam. It resides on a domain with e-mail addresses that is not used for any legitimate e-mail communication. The e-mail addresses from this domain are then posted on websites and other places from which spammers typically harvest e-mail addresses, when compiling their distribution lists. When the spammers eventually send out spam to these addresses, the spam trap logs the spam message, the time it is received, the IP addresses of the spam sending machine, and other details of the connection.

We then used this data to identify the location of the infection machines that sent the message – both its geographical location, as well as the network it belongs to. The infected machines will belong to a variety of different botnets. Many botnets are used to distribute spam. Of course, the spam trap will capture only a portion (or sample) of all infected machines existing worldwide, as not all bots will end up sending spam to any one particular trap. Given the gigantic volume of global spam, even a sample can become relatively large. The spam trap used as the basis of this report logged around 109 billion messages and 170 million unique IP addresses in the period 2005-2009. When compared to the spam data published by security companies, our data is congruent with their reports (see Appendix 1).

This resulted in data on just over 200 ISPs (see Appendix 1). Together, these ISPs control the bulk of the market share in the 40 countries. To cross check the completeness of our market data (as drawn from the Telegeography GlobalComms database), we compared it to the publicly available data on the total number of Internet subscriptions in each country (see Appendix 3). These public sources of data have their own shortcomings, as they rely on reports by the countries themselves, not on direct measurements. Still, if we use the 2009 OECD broadband statistics as a base of comparison, then the ISPs in our analysis account for 89 % of the total market in the OECD.

The process of mapping ASNs to ISPs was done manually. First, using the GeoIP data, we could identify which ASNs were located in each of the 40 countries. ASNs with 1 % of their IP addresses mapped to one of the 40 countries were included in our analysis. For each of these countries, we listed all ASNs that were above a threshold of 0.5 % of total spam volume for that country.

We used historical WHOIS records to look up the name of the entity that administers each ASN in a country. We then consulted a variety of sources – such as industry reports, market analyses and news media – to see which, if any, of the ISPs in the country it matches. In many cases, the mapping was straightforward. In other cases, additional information was needed – for example, in case of ASNs named after an ISP that had since been acquired by another ISP. In those cases, we mapped the ASN to its current owner.[17]

Figure 1. Number and location of infected machines over time
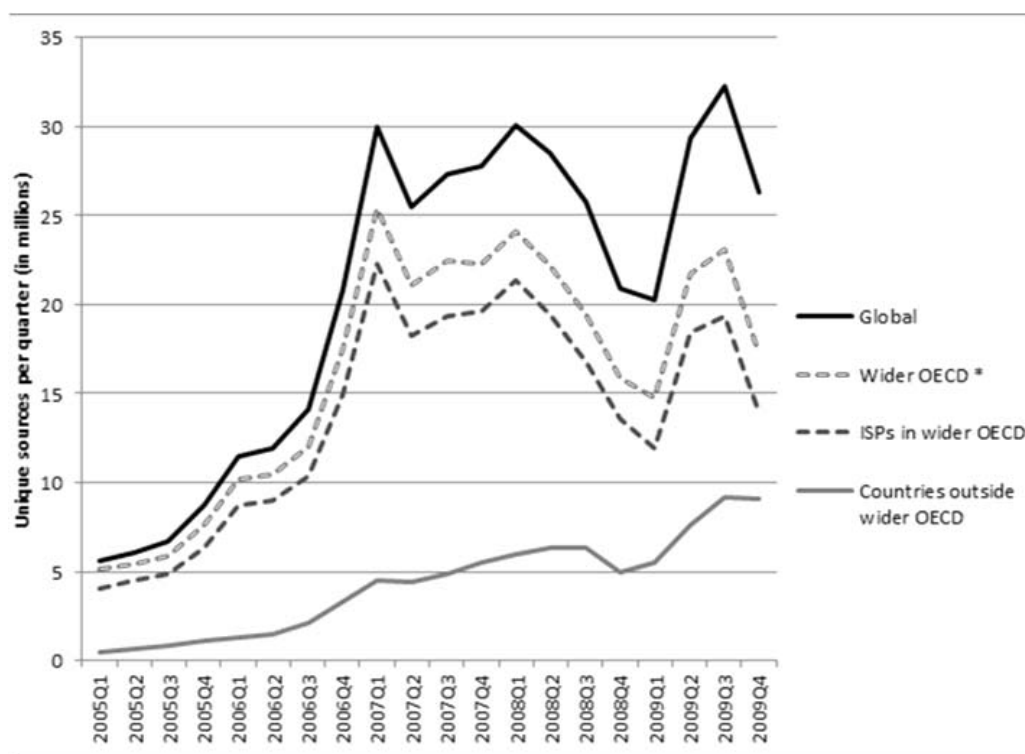
While we believe this to be a robust approach to establish the number of infected machines in ISP networks, it has certain limitations. Certain technical practices of network operators need to be take in into account when interpreting the data, most notably the use of Network Address Translation (NAT), the use of dynamic IP addresses with short lease times, and the use of port 25 blocking. In Appendix 4, we discuss how we compensate for the effects of these practices on the data. In short, we employ three different ways of measuring the number of infected machines in a network: the daily average number of unique IP addresses sending spam in a network, the yearly total number of unique IP addresses sending spam in a network, and the volume of spam from a network per year. All patterns that are discussed in the report are checked against all three measurement strategies. For the sake of brevity, we focus our discussion on the average daily number of unique sources. When spam volume or the total number of unique sources per year show a different pattern, we explicitly include it in the discussion. Where they are not mentioned, they are consistent with the findings as reported here.

The result of this approach is time series data on the number and the location of infected machines across countries and ISPs (Figure 1 shows one generic representation of that data). We have paid special attention to whether these machines are located inside or outside the list of countries covered by this study and to the degree in which they belong to the networks of the main ISPs in this geographic area. In general, the data reveals a rising trend in the number of infected machines, though the pattern is volatile, partially because of the arms race among attackers and defenders. With this data in hand, we can now turn to answering the research questions.

23

ARE ISPs CRITICAL CONTROL POINTS?

The most important reason to focus on ISPs as intermediaries is that they are asumed to be critical control points. As we discussed earlier, a related but usually implicit assumption is that the established ISPs – *i.e.* the known brands in the market – are important control points for infected machines, not the smaller, often fleeting players in the margins of the market.

As far as we know, these assumptions have never been emprically tested. Our data allows us to do just that. As explained above, we are working with a set of over 200 ISPs in 40 countries – 33 OECD member countries, two OECD Accession countries and five OECD Enhanced Engagement countries. This set consists of the ISPs that collectively posses the bulk of the market share in these countries. We first looked at the portion of the total number of unique sources of spam that can be attributed to these ISPs.

Over the period 2005-2009, between 60-74 % of all infected machines – *i.e.* spam-sending IP addresses worldwilde – were located within networks of around 200 ISPs in the wider OECD area. For spam volume, the numbers are slightly lower: 53-65 % (see Figure 2). If we look at the sources located in the 40 countries themselves, then that ratio is, of course, even higher: 80-83 % of all sources in this region reside in ISP networks.

**Box 3. Bulk of all infected machines are located in the networks of well-known ISPs**

As far as we can tell, all ISPs harbor infected machines – 'bots' – in their networks.What is surprising, however, is that the bulk of the total global population of infected machines are located in the networks of well-established providers, the brand names that are familiar to the consumers in those countries. Of the tens of thousands of ISPs that provide Internet access, the 200 ISPs that collectively hold nearly 90 percent of the total market share in the wider OECD area account for more than 60 percent of all infected machines worldwide. Other service providers, such as hosting providers, university networks, corporate networks and application service providers contain a smaller share of all bots.

This finding confirms the first assumption, namely that the bulk of infected machines worldwide is located in the networks of the established, predominantly retail ISPs in the wider OECD area – rather than in the networks of hosting providers, webmail providers, large corporations, universities, or application service providers. The pattern holds over time, although there seems to be a downward trend, from 74 % in 2005 to 60 % in 2009. The number of infected machines in networks of the 200 ISPs has gone down. More importantly, there has been an increase in the number of infections outside the wider OECD area. It is unclear to what extent this constitues a real trend. The underlying pattern is much more volatile than the yearly totals suggest (see Figure 1).

It is interesting to note that the portion of infected machines that are located in ISP networks varies significantly across countries. On the high end, we have countries like Israel, Turkey and Italy, where in 2009 over 90 %of all sources are located with ISPs. On the low end, we find Canada, with around 47 %, which might be explained by the fact that Canada has a large hosting provider industry that contributes a significant number of spam sources.

**Figure 2. What portion of all infected machines is located in ISP networks?**

Percentage of sources/volume of 200 ISPs in wider OECD compared to global total

---
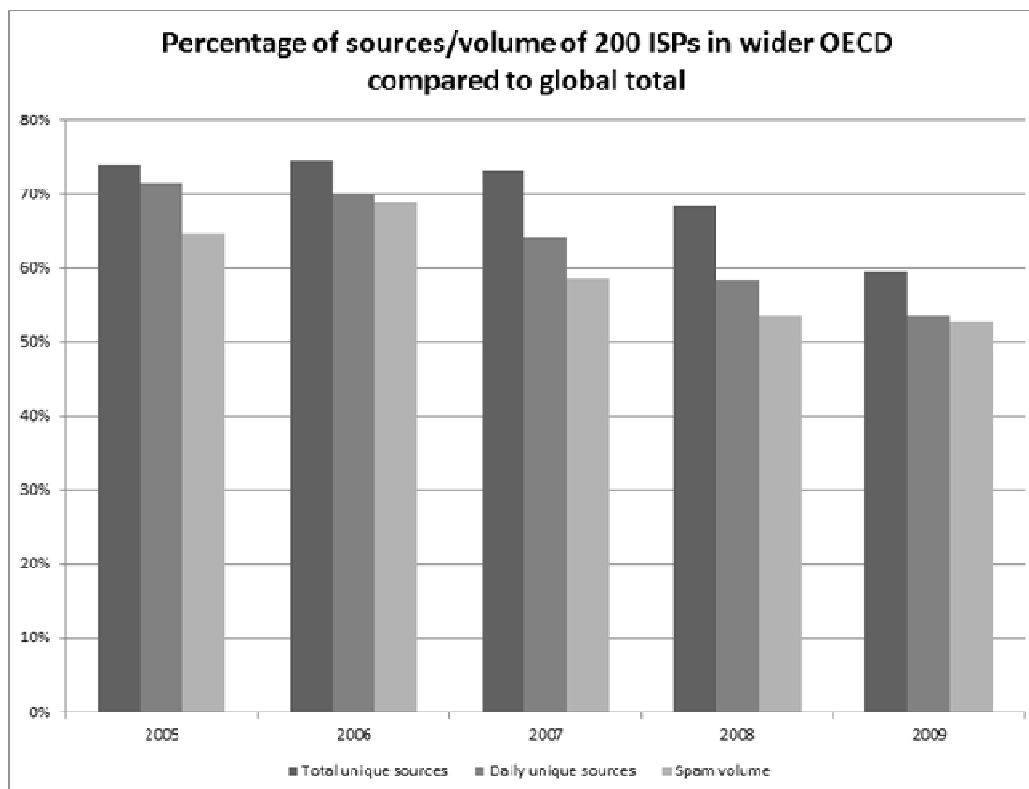
Box 4. Are Botnets migrating?

There is widespread belief, couched in experience, that any succesful security strategy will be countered by the attackers. The result is that the security problem, rather than being eliminated, often migrates elsewhere, to areas that are less well defended. This line of thinking would predict that if ISPs in the wider OECD area better defend their networks, then the botnets will compensate by migrating to other networks.

Many ISPs in the wider OECD area did in fact increase their efforts to combat botnets. And Figure 2 suggests a downward trend: a smaller portion of all bots are in their networks (from 74 % to 60 %). Does this mean that botnets are migrating?

There is indeed a slight decrease in the absolute number of infected machines in ISP networks. But it is unclear whether this constitutes a trend. Figure 1 shows how volatile the underlying pattern is. A better explanation for why, in percentage terms, ISPs in the wider OECD area now contain a lower portion of all botnets, is that the number of bots has increased in countries outside the wider OECD area (see bottom trend line in Figure 1). The latter most likely reflects the faster growth of the number of Internet users and subscriptions in those countries. The better their infrastructure becomes, the more they contribute to the problem of botnets. This is not a matter of migration, but of a growing global problem.

25

Figure 3. Concentration of infected machines in top ISPs
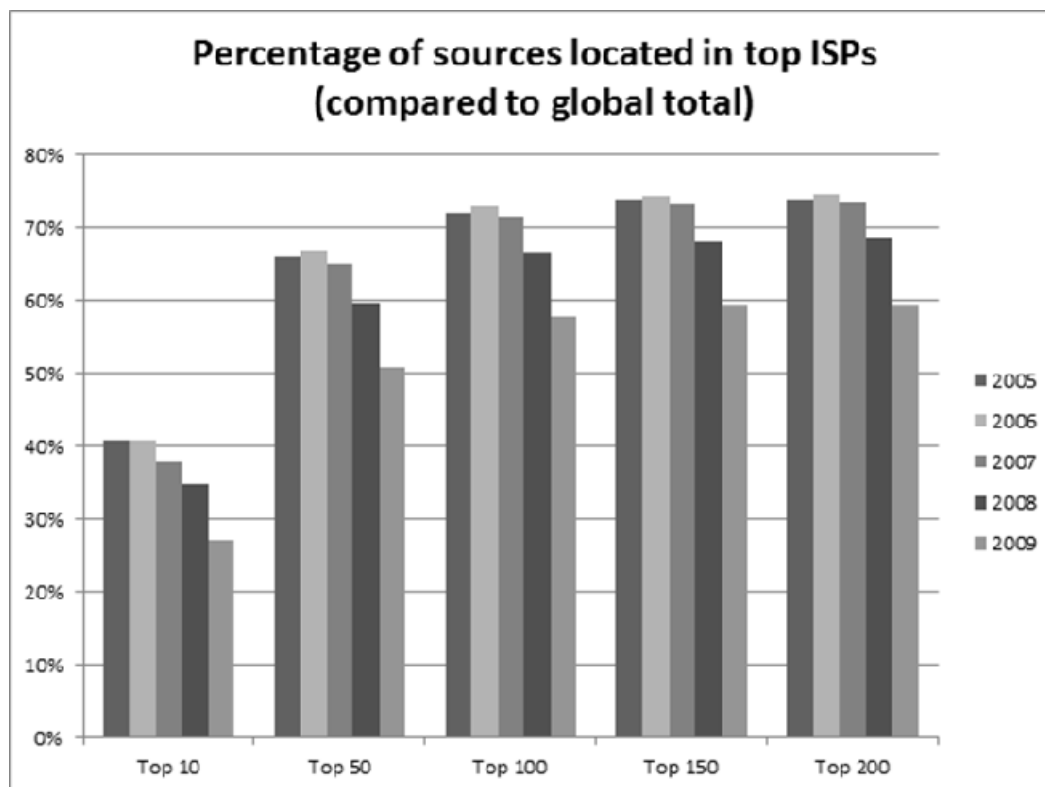
It is also intriguing to look at the distribution of infected machines within this set of around 200 ISPs. If we rank the ISPs by the number of infected machines in their networks in 2009, we find that the 10 highest ranking ISPs account for around 30 % of all sources worldwide (Figure 3). The top 50 ISPs account for around half of all sources worldwide. In light of the fact that there are 30 000 ASNs and anywhere between 4 000-100 000 ISPs, this is a remarkable finding. We also see that the curve flattens quickly. Adding the next 150 ISPs captures only an additional 7-9 percentage points of sources worldwide.

In light of the many thousands of players that are involved, collective action would seem an almost futile pursuit, given all the typical problems of free rider behaviour and weakest-link security. For botnet mitigation, however, the task of combating infected machines may have more manageable proportions, institutionally speaking. Our findings strongly suggest that the more established ISPs are indeed the ones which form critical control points, not the thousands of smaller players that it would be difficult, if not impossible, to reach through collaborative or regulatory efforts.

Of course, none of this is to say that improving botnet mitigation has suddenly become an easy task. Nor may such an approach offer a permanent reprieve as attack tactics could shift in respones to any mitigation measures, for example, by shifting more of the activity to the smaller players.

While we believe that the pattern we uncovered is relevant for the fight against botnets, we are not arguing that the same pattern holds across all activities in the botnet economy. Rogue providers play an important role in the command and control infrastructure of botnets, for example.[18] When the small hosting provider McColo was forced offline, it had a massive, if short-lived, impact on the global spam volume. This is not because the botnet machines themselves were located at McColo, but the machines that give the

bots their instructions. In those parts of the underground economy, criminal activity does, in fact, thrive because of weakest-link problems among ISPs. It only takes one hosting provider to enable command and control over one of more botnets. But these parts of the botnet economy would be useless without the actual bots, the infected machines. This is why it matters that the infected machines are predominantly located with a relatively small number of established providers. We revisit this discussion in more detail later, see Box 14.

Box 5. Even good ISPs tackle only a fraction of the bots

In a five year period, we found around 170 million different infected machines in our data alone. One could counter that this number would be lower when we take dynamic IP addresses into account. However, we only captured a sample of all spam sources, so in all likelihood the number is too low, rather than too high. This is confirmed when we consult other data sources on botnets, sources not based on spam. For example, we analysed data from the Conficker Working Group, which has logged the IP addresses of many machines that were part of the Conficker botnet. In 2009 alone, the Working Group logged around 169 million unique IP addresses. When we compared those addresses to the ones in our set of infected machines, we found that there is hardly any overlap among these sets. This means that each data source on its own greatly undercounts the total number of infected machines.

If we use the conservative estimate of 170 million infected machines, then this allows us to get a sense of the effort required by ISPs if they would really contact all infected customers. We use an average-performing, medium sized ISPs of just over 2 million customers as an illustration. The total number of infected machines detected in 2009 in the network of this ISP, divided by 365, suggests that the company would have to contact around 100 infected customers per day. The actual number of infected machines we see active every day is much higher, averaging over 750. But many of those show up on multiple days and we count them only once. So contacting 100 customers is the minimum effort, assuming all customers immediately clean up their machine after a single notification.

There is no data available about how many infected customers are being contacted by ISPs every day. But the anecdotal evidence we have gathered suggests that the current efforts are one order of magnitude below the actual number of infections – that is, the ISP in our example contacts around 10 customers every day, and we suspect that it is among the more active ISPs in this respect.

DO ISPs PERFORM DIFFERENTLY IN TERMS OF BOTNET MITIGATION?

A lot has been written about the incentives of ISPs, or lack thereof, to improve security. Various incentives have been identified, some enhancing security, others working against it. It is not at all clear what the net effect is of these incentives on ISP's behaviour, nor whether this effect varies significantly across ISPs. Another way to frame this problem is to ask how much discretion ISPs have in mitigating botnets. If they are subject to similar, externally imposed, incentives that leave them little organisational freedom to respond to them, then we would expect similar performance in this area. However, if they do have discretion and can respond differently, diverse performance outcomes will be observable.

In order to compare the number of infected machines across ISPs, we need to take into account a number of factors. The first factor is size. Our dataset includes a range of ISPs of varying size. In 2005, the smallest ISP had 3 000 customers and the largest 21 million customers, with the median at approximately 247 000. In 2009 the numbers were higher, with the smallest ISP showing 13 300 and the largest 53.5 million customers. The median in 2009 was at approximately 560 000 customers.

Obviously, other things being equal, ISPs with more customers will experience more infections. If we look at ISP performance – as measured by the number of infected machines – and rank the ISPs according to their size, this becomes immediately visible (see Figure 4 for the 2009 findings). When the number of subscribers and the number of infected machines are both logarithmically transformed, we can see a nearly linear relationship: infection rates rise together with the number of subscribers.


Box 6. Number of customers is the strongest factor influencing the infection rate of an ISP

The group of ISPs in our analysis make up a highly varied group, not only in terms of security policies, but also in terms geography, market position, the characteristics of users they serve, and the regulatory frameworks under which they operate. All of these factors potentially influence the number of infected machines found in their networks.

Notwithstanding all this variety, however, our analyses clearly show that the number of customers is the largest contributor to the number of infections. The number of customers alone, leaving out all other factors, can explain about 43% of the variance in the infection rates we measured across the ISP networks. In other words, more customers means more infections, no matter what ISP policies are in place, how technically competent its users are or how active the government is in promoting cybersecurity. This implies that we are dealing with a systemic problem that affects all ISPs, not with a problem that can be reduced to an issue of good versus bad ISPs.

The size of the ISP can explain the infection rates in these networks to some extent – in statistical terms: size explains 43 % of the variance for the period 2005-2009. That being said, we can also see a remarkable degree of variability. Across the board, there is a difference of two orders of magnitude, sometimes even higher, in the number of infected machines within networks of ISPs of similar sizes. In other words, some ISPs harbour about a hundred times more infected machines in their network than their peers of similar size. This is not a matter of outliers.[19]

Figure 4. Infected sources and number of subscribers of ISPs in the wider OECD area (2009)
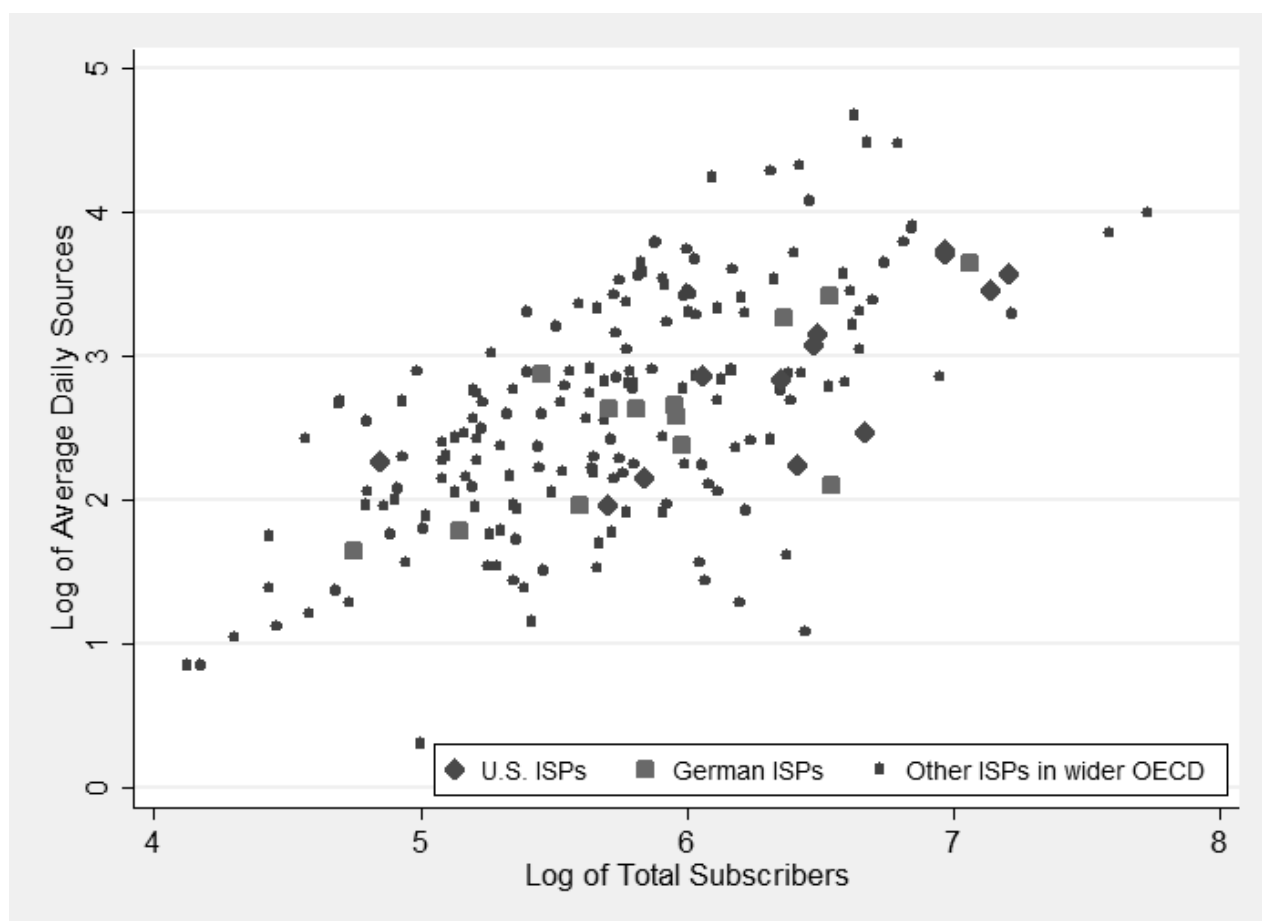
Other factors may be country-specific, such as legal framework or the cost of customer support. But even within countries, where we can assume that ISPs operate under similar institutional incentives, we see substantial differences in performance. In the United States and Germany, for example, we still see at least one order of magnitude difference, often more, among ISPs of similar size (Figure 4).

A third aspect is the performance of ISPs over time. Although we did not perform a detailed empirical analysis of the dynamic response of individual ISPs to infections on their networks, we conducted an aggregate analysis of ISPs in the set of worst performers. In each year, we looked at the 50 worst performers among the ISPs that we studied – roughly containing about half of all infected machines in their networks. How dynamic is the composition of this group of poor performers? Is ISP performance volatile and dominated by the strategies of attackers?


Box 7. Similar ISPs can have a tenfold difference in infection rates

When we compare ISPs that operate under very similar conditions, we see remarkable differences in the infection rates. ISPs of similar size, operating in the same market and under the same regulatory framework, show infection rates that vary by a factor of ten or more. That is, one ISP has ten times more infected machines in its network than its very similar peer. This seems to suggest that factors at the ISP level, such as its security policies, make a significant difference. To put it differently, external conditions do not fully dictate how much ISPs can do about botnet infections. They have considerable leeway to mitigate botnets.

The Venn diagrams in Figure 5 illustrate overlaps in the 50 worst performing ISPs between 2006 and 2009, both for the total number of infected machines and for the number of infected machines per subscriber – *i.e.*, absolute and relative infection rates. There is some variation in membership in the total set and the various sub-sets. For example, a total of 71 ISPs were in the top 50 in one of the four years based on the number of infected machines daily, and 78 ISPs were in the top 50 based on the number of infected machines per subscriber. However, we also observe a stable core of 31 ISPs that had the highest number of infected machines on their network during all four years (13 ISPs were in the set in three years, 10 in two years, and 17 in only one year). With regard to infected machines per subscriber, 26 ISPs were in the top 50 during all four years, 15 were in it during three years, 14 during two, and 23 during only one year.

These observations suggest that there is considerable inertia at the core. Statistically, the performance of an ISP can be predicted to a large extent – 68 %, to be precise – from the performance of the previous year. Notwithstanding the volatility caused by changing attacks, it seems ISPs respond predictably to this volatility. That being said, some ISPs apparently were able to improve their performance and hence moved out of the group. However, others struggled with deteriorating performance and hence moved in, possibly temporarily.

Figure 5. Variability in top 50 ISPs with the most infected machines in 2006-2009

(a) Variability in total number of infected machines

(b) Variability in number of infected machines per subscriber
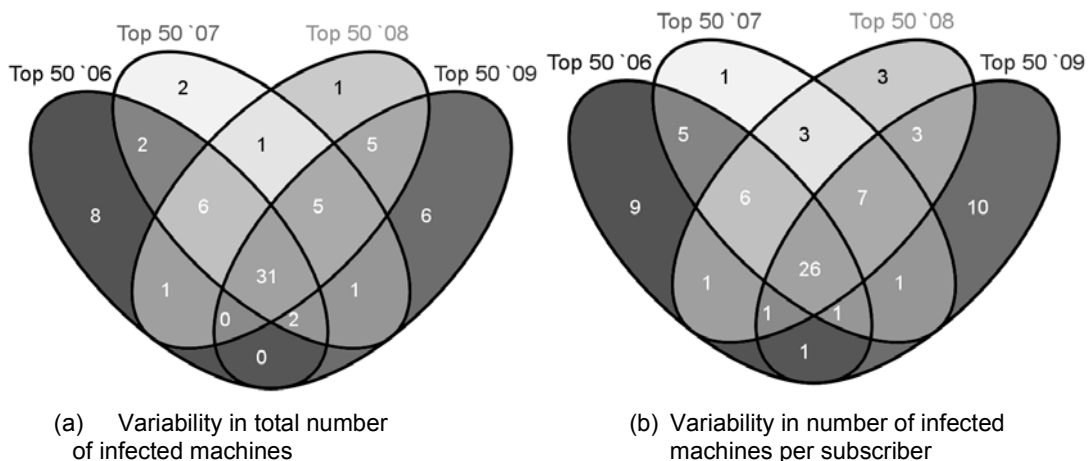
Box 8. Most-infected ISPs are distributed cross the wider OECD area and include ISPs of all sizes

What ISPs are among the core set of most infected networks? As has been discussed earlier, we identified 31 ISPs that consistently contributed the highest number of infected machines over the past five years. We do not think it is appropriate to list the names of those ISPs. More research is needed to robustly benchmark ISP performance. (As an aside, we are currently undertaking such research for the Dutch market, a study commissioned by the Dutch government.)

Without mentioning names, we can however indicate in more general terms what ISPs are part of the group of poorest performers. Geographically, they are distributed across 17 countries in the wider OECD area (Figure 6). We know that size of the ISP – in terms of the number of customers – is an important factor that influences the number of infected machines in their network. But these 31 poorerst-performers are not just the largest ISPs. They range in size from 197 000 to 53.5 million subscribers (Figure 7, subscription numbers for 2009). The lower end of that range includes relatively small ISPs, well below the median (552 000 subscribers) and the average (1.9 million subscribers).

When we look at relative infection rates – that is, the number of infections per customer – we found a stable set of 26 poor-performing ISPs. They are also distributed georgraphically, in this case across 12 countries. The range of sizes extends more to the lower end of the ISP population, from 11 900 to 6.2 milllion subscribers. This is consistent with findings that we discuss later in the report, namely that large ISPs, on average, have fewer infections per subscriber.

Figure 6. Geographical distribution of the most infected ISPs

Figure 7. Size of consistently most Infected ISPs (2009)
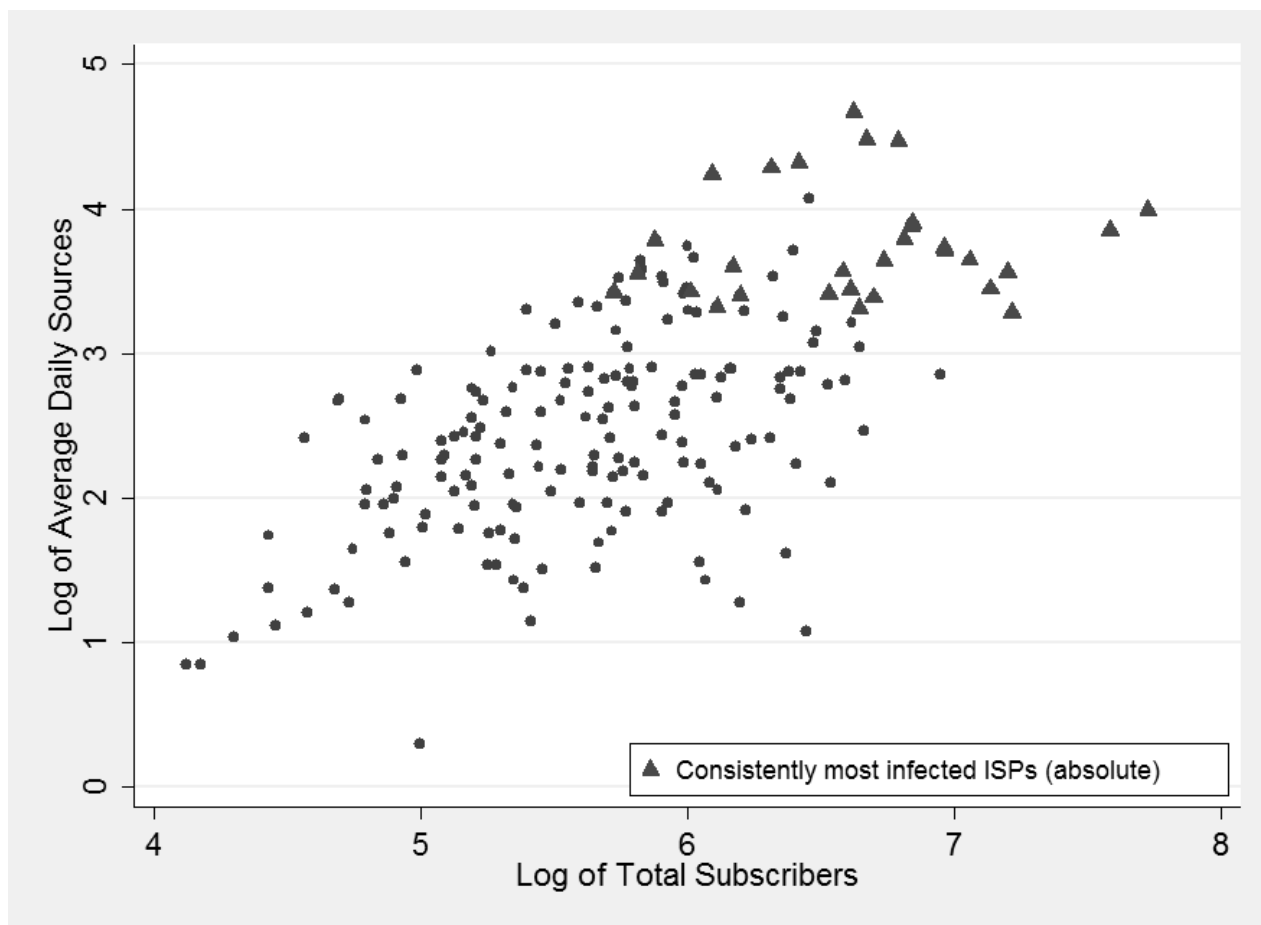
So the group of most infected networks is relatively stable over time. Who are the ISPs that administer these networks? Geographically, they are distributed across a variety of countries (Figure 6). In terms of size, they are also quite diverse. In the group of 31 networks that contain the highest number of infected machines, the smallest ISP has 197 000 subscribers, the largest one 53.5 million. The group of 26 networks that have the highest relative infection rates, *i.e.* the most infected machines per subscriber, the smallest ISP reported about 11 900 subscribers and the largest one around 6.2 million subscribers. Again, the finding is reinforced that these infected networks are established brands, not rogue players hiding in the margins of the market.

In sum: All of this suggests that ISPs have substantial discretion to decide how they engage in botnet mitigation and that their organisational incentives lead to different choices, even when working under a common set of institutional incentives, such as defined by the legal framework of a country. This point is reinforced when we look at the differences between countries, rather than ISPs. At the country level, our data measures the total spam output of ISPs and non-ISPs. As players with very different records within one country are aggregated, country performance data show less variance than individual organisation data. Consequently, at that level of analysis, the number of Internet users explains around 70 %of the variance in performance. As ISPs do perform very differently under comparable institutional incentives and economic circumstances, this suggests that country-level mitigation measures, while necessary, will not be sufficient unless they also address the organisational incentives and realign both. In the next section, we explore the extent to which we can explain these differences among ISPs.

EXPLAINING THE DIFFERENCES AMONG ISPs

Conceptual framework

Advanced information and communication technologies form a highly interrelated ecosystem. Like other actors, ISPs respond to economic and non-economic incentives. Most generally speaking, incentives are the factors that individual and organisational decision-makers take into account. Given the highly dynamic nature of this ecosystem, the observations reported in the previous section could be the emergent outcomes of the varied responses by ISPs to the problems of botnets without an underlying stable pattern. However, if the phenomenon had certain regularities this knowledge could be utilised to improve cybersecurity. We therefore formulated a simplified conceptual model of the ecosystem around ISPs and subjected it to empirical analysis. Figure 8 represents a stylised model of the factors that influence botnet activity: the security measures adopted by an ISP, the level and virility of cybercriminal attacks, technological factors, and user behaviour. Other factors, such as the behaviour of software vendors and registrars, also impact this ecosystem, but they are outside the scope of this study (see van Eeten and Bauer 2008 for an in-depth discussion). An ISP's decision to adopt security measures is influenced by factors related to the institutional and organisational environments. These groups of factors are linked in multiple feedbacks so that they co-evolve over time. For example, stronger security efforts by an ISP may reduce botnet activity but also result in stronger efforts by cybercriminals to find new attack vectors. As our units of analysis are ISPs, it is important to take the national context into account. However, cybercrime is a transborder phenomenon and the international context is therefore also relevant.

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

The level of effort that ISPs exert on botnet mitigation depends on the relevant set and the relative strength of positive and negative incentives. Relevant institutional factors include the legal and regulatory framework in which ISPs operate, the market structure and the associated competitive pressures, and the conditions in related markets, *e.g.* for security technology. Relevant organisational factors include the size of the customer base, the organisation of the abuse desk, and the cost of various security measures. Which incentives will be perceived as relevant by an ISP is influenced by its business model. Commercial ISPs will primarily respond to incentives that have direct and indirect implications for their bottom line.

Likewise, rogue ISPs deriving most of their business from activities related to cyberfraud and cybercrime will also primarily respond to economic incentives (Bauer and Van Eeten, 2009). In both cases, non-economic incentives, such as peer pressure and peer recognition, may play a role. These types of incentives are often seen to be subordinate to economic incentives. This need not be the case, however. When peer pressure takes the form of blacklisting, it has economic effects that can be quite significant, such as rising cost of customer support, when customers experience the effects of blacklisting and start calling their ISP. The relative weights of relevant incentives could be different for non-profit ISPs or co-operatives but even such ISPs do not have unlimited resources and will have to pay attention to economic factors. All ISPs will therefore be influenced by the incentives identified in Figure 8, which interact to jointly influence an ISP's botnet mitigation effort.

The signs in parentheses in Figure 8 refer to the direction of the incentive, other things being equal. A positive sign indicates that the incentive works in the same direction as the factor next to which it is listed. For example, tougher cybersecurity laws will likely increase the level of botnet mitigation by an ISP. The same holds in the opposite direction: weaker laws will coincide with weaker incentives. Thus, the factor and the incentive work in the same direction, resulting in a positive sign. A negative sign indicates that the factor and the incentive move in the opposite direction. For example, the presence of a business model built around malicious activity will reduce the incentives to invest in security (thus a negative sign). The strength of an incentive is quite a different issue and may depend on the presence of complementary incentives. For example, laws providing a base for action against spammers will be more effective if they are also enforced actively. Likewise, the effectiveness of liability rules, which are often mentioned as a possible course of action, will depend on whether or not the required burden of proof can be met. Because this encounters great difficulties, leading legal scholars tend to be skeptical whether liability rules are workable (*e.g.* Spindler 2007). The effectiveness of incentives and the interaction between them will also be influenced by the national context. The diffusion of broadband service, the income level of a country, the education level of the population, and the diligence of law enforcement might be of particular importance. However, the actual strength and directionality cannot be determined on the basis of theoretical reasoning alone but needs to be supplemented by empirical analysis.

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

Figure 8. Conceptual framework
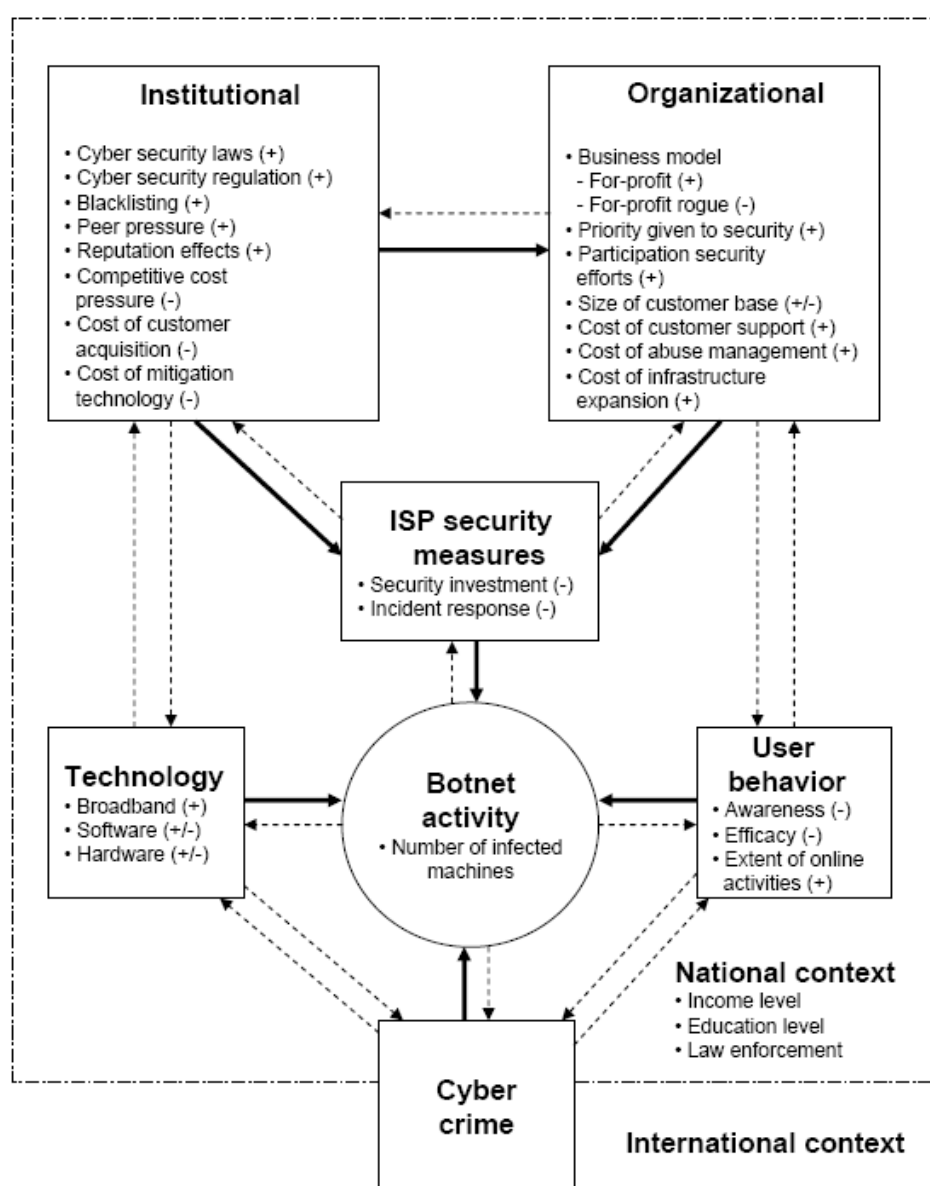
For purposes of empirical analysis, data constraints necessitated further simplification of the theoretical approach to a more manageable empirical estimation framework (Figure 6). We used two proxies to measure the number of infected machines: the number of unique IP addresses emitting spam and the total number of spam messages originating from an ISP during a specific time period. Drawing on the conceptual framework discussed above, a large number of variables that could either serve as direct measures of proxies for the independent variable were used. In this paper, only variables that were seen as relevant based on the conceptual approach and that yielded a statistical contribution are reported. In addition, control variables were introduced to take factors related to technology, user behaviour, and the national context into account. As cybercrime is a globally mobile phenomenon, we proceeded on the assumption that all ISPs are targeted at a comparable rate (although the level of botnet activity is influenced by an ISP's security efforts as well as other control factors and will thus vary). The empirical model is displayed in Figure 9.

Figure 9. Empirical framework
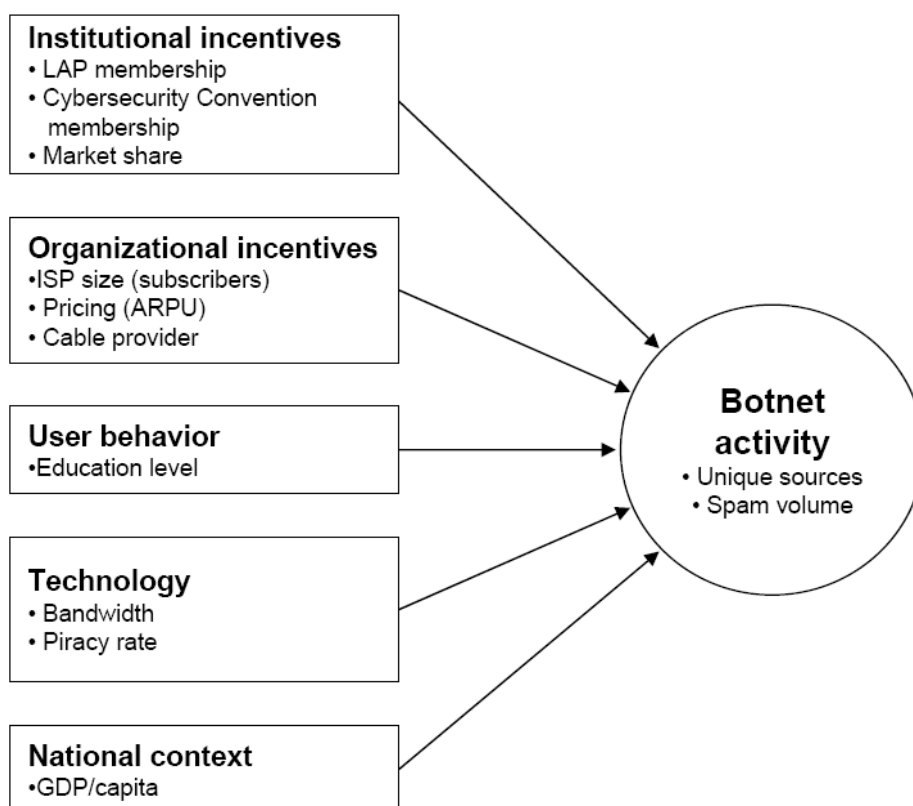
Data for the independent and control variables was collected from several sources, including the World Bank's World Development Index database, the UN Human Development Reports, the Business Software Alliance, and TeleGeography's GlobalComms database (see Appendix 2). Where possible, data was triangulated against other sources, such as the International Telecommunication Union's World Telecommunications Indicator database.

In addition to the 109 billion spam messages from 170 million unique sources which were parsed, aggregated, and attributed to ISPs and countries in the way discussed above, we were able to assemble a panel of annual observations for 2005-2009 for 40 countries. Although we were able to gather considerable evidence, it was not possible to generate empirical data for all the variables suggested by the theoretical model forcing us to work with proxies where available. However, in some cases, such as prices for Internet access services, the empirical model was constrained by lack of data.

Empirical findings

The dynamic nature of botnets raises many methodological challenges and complicates their statistical analysis. To gain different insights into the empirical relations, three approaches, each with its own advantages and disadvantages, were used to examine research questions derived from the theoretical framework: *i)* bivariate methods, *ii)* multivariate methods using a pooled data design, and *iii)* multivariate methods using panel data analysis. In each case, we set out to explain the differences in botnet infection rates among the ISPs. Empirical analyses were conducted for multiple dependent variables, including the total number of spam messages emitted from an ISP during a specific time period, spam messages per

subscriber, the number of unique IP sources emitting spam from an ISP during a specific time period, and the number of infections per subscriber (see Appendix 2 for a complete list). Results presented here use the most reliable metric, the daily average number of unique sources per subscriber.[20] The findings we report are consistent with those for the other dependent variables, unless stated otherwise.

Bivariate analysis offers a first crude look onto the relations between independent variables and the proxies for botnet activity. However, it has to be kept in mind that such simple correlation analysis neglects the effects of other factors that may play a role and therefore may attribute too much influence to a single independent variable. The approach allows only a preliminary, descriptive understanding of the data structures. Therefore, the robustness of findings has to be ascertained using other methods. Depending on the nature of the data, the methods for testing the association of an independent variable and daily average unique sources per subscriber varied, including rank correlation coefficients and t-tests. The results of this analysis are presented in Table 1.

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

- There is a widely held belief that larger ISPs show worse security performance, as they face much less peer pressure. For instance, Moore, Clayton, and Anderson (2009) state that "...very large ISPs are effectively exempt from peer pressure as others cannot afford to cut them off. Much of the world's bad traffic comes from the networks of these 'too big to block' providers." In contrast to this belief, our dataset indicates that, while larger ISPs emit more spam in absolute numbers, relative to size their performance is on average slightly better than that of smaller ISPs.

- Another claim is that lower average revenue per user (ARPU) is a sign of higher financial pressure that might result in less attention to security. Our data suggests that ARPU and relative security performance are unrelated.

- Given differences in networking technology and user base, one might hypothesise that cable service providers can enhance their security performance easier than DSL providers. Our data indicates an 8 % lower incidence of unique sources for cable companies. The volume of spam, however, is similar for both types of providers. This might reflect that cable subscriptions have higher average bandwidths than DSL subscriptions, that cable providers use more Network Address Translation technology, or that they more often block port 25.

- Bivariate analysis indicates that ISPs in countries that have joined the London Action Plan (LAP) have, on average, fewer bot infections. Likewise, operating in a country that has signed the Council of Europe's Convention on Cybercrime is negatively correlated with botnet infections. Neither of these initiatives targets botnets directly. However, one could argue that membership of LAP is a proxy for the activity of a country's regulatory entities in the area of cybersecurity, whereas membership of the Convention on Cybercrime is a proxy for the activity of law enforcement institutions in a country. These memberships, we assume, are associated with a broader set of measures undertaken by the governments in those countries. Earlier research by Wang and Kim (2009) provided some evidence in support of this effect, though they presume a somewhat tenuous direct causal link between the Convention and cybercrime incidents, rather than interpreting membership of the Convention as a proxy variable. However, factors correlated with a country's willingness to sign these agreements could also be at work both for the Convention as well as the LAP.

Table 1. Correlation analysis of the factors influencing the number of infected machines per subscriber*

Subject	Independent variables	Statistical instrument	N	Results *pooled, daily avg*	Results *pooled, all sources*	Results *pooled, volume*	Results
Effects of ISP size	total_sub	Spearman's rho	932	$\rho = -0.225$	$\rho = -0.157$	$\rho = -0.182$	**Negative relation**
	market_share	Spearman's rho	699	*insignificant*	*insignificant*	$\rho = -0.079$ (sig=0.038)	**No relation**
Effects of ARPU	rev_persub	Spearman's rho	148	*insignificant*	*insignificant*	*insignificant*	**No relation**
Cable vs. DSL providers	srv_cable	t-test	828	*diff = .0005*	*diff = .0975*	*insignificant*	**Cable providers have fewer sources**
Effects of regulation	lap_mem	t-test	932	*diff = .0013*	*diff = .120*	*diff = 48.7*	**LAP members have fewer sources**
	cyber_mem	t-test, K-Wallis	932	*diff = .0016*	*diff = .130*	*diff = 51.4*	**CC members have fewer sources**
Effects of piracy	piracy_rate	Spearman's rho	930	$\rho = 0.506$	$\rho = 0.436$	$\rho = 0.383$	**Positive relation**
Effects of bandwidth	int_bpp	Spearman's rho	383	$\rho = -0.209$	$\rho = -0.227$	*insignificant*	**No relation**
Effects of user educ.	educ_ix	Spearman's rho	932	$\rho = -0.408$	$\rho = -0.412$	$\rho = -0.278$	**Negative relation**

* The bivariate statistical tools used are comparison of means and measures of association. In this table, unless stated otherwise, all reported test results are statistically significant at the 0.01 level; ρ is the rank correlation coefficient, and is a measure of the degree of association of two variables (between -1 to 1); *diff* is the difference between the averages of the two sample groups. Due to lack of normality in the dependent variable, non-parametric tests were often employed.

- A frequently stated claim is that countries with higher rates of software piracy also have higher botnet activity. We have tested this claim using data from the annual BSA/IDC Global Software Piracy Study.[21] At the bivariate level, we found a moderate positive relation between piracy and infection rates.

- Bandwidth is often seen as enabler of malware (*e.g.* OECD 2009). However, our data does not support that claim at the bivariate level and we did not find an indication that increased use of broadband connections directly translates into a higher number of bot infections – measured either in the number of infected sources or spam volume.

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

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

The next step in multivariate analysis was to model the factors that influence the performance of ISPs. For this purpose, the number of botnet infections was corrected for an ISP's size. Such relative measures of botnet activity (average number of sources per subscriber or spam volume per subscriber) are more appropriate for a comparison of ISPs. Within the constraints of data availability, a broad range of factors at the country-level (*e.g.* a country's LAP membership status) and the ISP-level (*e.g.* subscriber base of the ISP) were examined. Moreover, we tested control variables reflecting the technical infrastructure and economic situation of a country.[2223] In a first round of analysis, we were seeking for a parsimonious explanatory model by assuming that institutional factors would affect infection rates in an additive fashion. In a second round we took into account that institutional arrangements frequently interact with each other and socio-demographic factors, such as the education level of a nation. *A priori* it was not clear whether the empirical relationship could be better characterised by a model with constant or with random coefficients. Therefore, we tested the influence of the explanatory variables under both assumptions (using pooled data for the constant coefficient approach and panel data for the random approach). Compared to the pooled model, the panel data approach allows extracting additional insights from the cross-sectional and time-series variation of the data. In other words the method takes advantage of the fact that data originated from different ISPs and was observed at different points in time.[24] This yields a total of four models (without and with interaction terms; fixed and random coefficients).

Table 2. Determinants of number of infected machines per subscriber

Explanatory variable	Pooled model		Panel model	
	Coefficient	Standard error	Coefficient	Standard error
No. of subscribers	-0.00456***	0.00112	-0.00198**	0.00081
Cable ISP	-0.00272**	0.00136	0.00050	0.00114
Cybercrime Conv.	-0.00055	0.00204	-0.00083	0.00165
LAP membership	-0.00735***	0.00168	-0.00807***	0.00142
Piracy rate	0.00041***	0.00007	0.00030***	0.00005
Education level	-0.05886***	0.01931	-0.06749***	0.01967
Constant	0.10528***	0.02181	0.09869***	0.02133
N	826		824	
Adjusted R^2	0.28		n.a.	
Joint significance	F = 50.22***		Wald = 239.49***	

Notes: Statistical significance at 1% (***), 5% (**), and 10% (*); n.a.: not available.

Representative findings from the models without interaction terms are presented in Table 2. The explanatory power of a model can be assessed by looking at the statistical significance of individual parameters and of the overall model. Moreover, the share of the total variance in the dependent variable that can be explained is relevant. The number of subscribers, LAP membership, piracy rate, and education level all show a statistically significant effect in the pooled and panel design. Status as a cable modem provider is only statistically significant in the pooled model. Moreover, as important, the sign of the parameters, which indicates the direction of the influence, is the same in both approaches. A larger number of subscribers are associated with a slightly lower average infection rate; countries that are LAP members have lower infection rates; and higher education levels go hand-in-hand with lower infection rates. On the other hand, ISPs in countries with a higher rate of pirated software show a higher average infection rate. Membership in the Cybercrime Convention, although associated with a negative sign, is not statistically significant in either approach. Both models are highly significant overall. The pooled model explains about 28 % of the variance in average sources per subscriber among the ISPs. Extended versions of these models, which we will discuss in a moment, explain about 40 % of the variance. That a large share of the variance remains unexplained indicates that there are other factors at work – in particular the dynamic behaviour of

attackers and defenders – that are not fully captured by the institutional and operational variables that were available for inclusion in the model.[25]

The findings from the pooled, constant coefficients model are largely consistent with the effects detected by simple bivariate analysis. One finding that is replicated is that as the subscriber base increases, the number of infections per subscriber decreases (as expressed by the negative sign of the coefficient for the log of subscribers). Simply put, large ISPs are, on average, doing slightly better than smaller ones. This contradicts the widely held view that large ISPs operate with impunity and achieve worse security performance than smaller ISPs.

Table 3. Determinants of number of infected machines per subscriber (with interaction terms)

Explanatory variable	Pooled model		Panel model	
	Coefficient	Standard error	Coefficient	Standard error
No. of subscribers	-0.00298**	0.00126	-0.00155*	0.00095
Cable ISP	-0.02781**	0.01253	-0.01252	0.00879
Cybercrime Convention	-0.00021	0.00212	-0.00188	0.00169
LAP membership	0.38135***	0.11361	0.11759	0.10072
Piracy rate	-0.00253*	0.00145	-0.00333**	0.00131
Education level	-0.28774***	0.09818	-0.34851***	0.08583
Joint effect cable and no. of subscribers	0.00415*	0.00217	0.00203	0.00153
Joint effect of LAP and education	-0.42645***	0.11867	-0.14384	-0.10474
Joint effect of LAP and piracy	-0.01246***	0.00185	-0.00665***	0.00169
Joint effect of education and piracy	0.00308**	0.00151	0.00378***	0.00136
Joint effect of LAP, education, and piracy	0.01380***	0.00199	0.00741***	0.00181
Constant	0.31564***	0.09575	0.36778***	0.08485
N	826		824	
Adjusted R^2	0.40		n.a.	
Joint significance	F=43.66**		Wald=403.60***	

Notes: Statistical significance at 1% (***), 5% (**), and 10% (*); n.a.: not available.

As institutional and other incentives typically do not work in isolation from each other, we also introduced interaction terms to examine whether the joint presence of selected factors was important. Given the lack of detailed ISP-level data, we could test for interaction effects predominantly at the country level. For example, we explored the influence of both LAP membership and high education levels on the average infection rates. Representative findings of models with interaction terms are summarised in Table 3.

Once such joint effects are allowed, a more complicated and multifaceted structure of effects is revealed. Nonetheless, the overall picture that emerged from the simpler analyses depicted in Table 2 is by

and large still visible. As in those more parsimonious models, the number of subscribers is negatively related to the average infection rate. Likewise, the findings replicate the earlier insight that cable providers have lower infection rates. If indeed, as we hypothesised earlier, this effect is tied to the use of automation, it can also explain the interaction term between subscriber count and cable access. The sign of the coefficient of the interaction term is 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 summing their individual effects (in other words, if an ISP already has automation in place due to being a cable provider, then size will not make a difference). When automation is present, it may cause lower infection rates because it drives down the cost of migration efforts, such as contacting, filtering and quarantining customers – a cost which is seen as an important incentive that works against botnet mitigation by ISPs.

In the pooled model specifications, the parameters of the membership in the Cybercrime Convention, and education levels of users all were negative, indicating that these factors mitigated botnet infection rates. However, compared to the bivariate analysis, the parameter sign of piracy rate switched to negative. Also, the parameter sign of LAP membership, taken by itself, changed to positive, suggesting that LAP membership alone is not a sufficient factor for lower infections. However, LAP membership jointly with other variables, such as education, continues to show a mitigating negative effect on infections. The negative parameter of the interaction term of LAP membership and piracy rate is somewhat less convincing (but could be explained by a self-selection effect in the sense that countries with higher piracy issues might also be inclined to be more active in LAP-related mitigation efforts). These changes compared to the simple correlation analysis are not unexpected and indicate that some of the findings are sensitive to the specification of the model and in that sense less robust than those findings that do not change (*e.g.* the negative effect of ISP size, the mitigating effects of government measures).

Box 9. Large ISPs have fewer infected machines per customer than small ISPs

Previous research found that peer pressure among ISPs is an important incentive that contributes to security. Some authors have claimed that large ISPs experience less peer pressure than smaller ISPs, because other ISPs cannot afford to cut them off. This lack of peer pressure would lead to poorer security performance, according to these authors. Our data, however, reveals that this claim is incorrect. Large ISPs have on average fewer infections per customer than smaller ISPs. This is not to say that smaller ISPs always perform worse. Many of them perform very well. But the variability among them is large and on average, they do worse than large ISPs. Elsewhere we discuss some evidence that suggests the better performance of large ISPs is related to the use of automation in handling the abuse issues of its customers. Automation reduces the cost of dealing with infections, which may allow the ISP to increase its efforts without raising cost.

In the panel specification with interaction effects, all variables show the same sign as in the pooled data model. However, the significance level of several variables differs compared to the pooled model. In the panel data several variables (total subscribers, cable, cable interacted with subscribers, LAP membership, Cybercrime convention membership, and LAP interacted with education) are not statistically significant whereas in the pooled model only the Cybercrime Convention membership is insignificant. Statistical significance may not be a central concern though, because the ISPs in our dataset represent the lion's share of the ISP market in the 40 nations – compared to the 2009 OECD broadband statistics, the ISPs in our analysis cover 87 % of the total market in the wider OECD area. Therefore, our data is nearly a complete enumeration of the markets rather than a sample. In this case, the parameter estimates reflect the data structures in the empirical universe under investigation. So far as ISPs are concerned, it is not necessary to make inferences from a subset to the whole phenomenon. Consequently, significance levels are less critical when interpreting the findings.

**Box 10. Governmental efforts seem to reduce infection rates at ISPs**

Our analyses uncover mixed evidence regarding the role of governments. However, the overall conclusion of the different findings seems to be that government efforts have a weak negative impact on botnets – that is, higher efforts lead to lower infection rates in their countries. When we look at the data more qualitatively, we find further support for this conclusion. Some countries are known for their efforts in which governments and ISPs collaborate in the area of cybersecurity, most notably Japan and Finland. More recently, Australia has also launched an initiative. In our data, Japan and Finland consistently show up among the countries with the lowest infection rates.

When we plot the number of infected machines in the ISPs networks in each country against the total number of subscribers, the countries that perform well can be found on the bottom of the cloud (Figure 10). Indeed, we see Japan and Finland there, as well as Australia and several other countries.

The same data can be represented differently, by calculating the average infection rates per subscriber of all ISPs in a country (Figure 11). Again, we find that the ISPs in Japan and Finland are, on average, among the least infected networks in the wider OECD area. To reiterate: on the whole, we cannot draw robust conclusions from our evidence, but there is certainly support for the finding that governmental efforts are associated with lower botnet infection rates at ISPs.


Designing a policy approach needs to take the relative effects of changes in instruments into account. Although individual parameters across the different modeling frameworks show largely congruent signs and statistical significance, their sizes differ somewhat. Our models do not suggest that the independent variables can be used as "levers" to mitigate botnets. However, they do suggest that ISP size matters, that government policies matter, and that the interaction of government policies and socio-demographic factors also matters. Table 4 tries to summarise at a high level these overall insights.

Figure 10.   Number of infected machines of all ISPs in each country (2009)
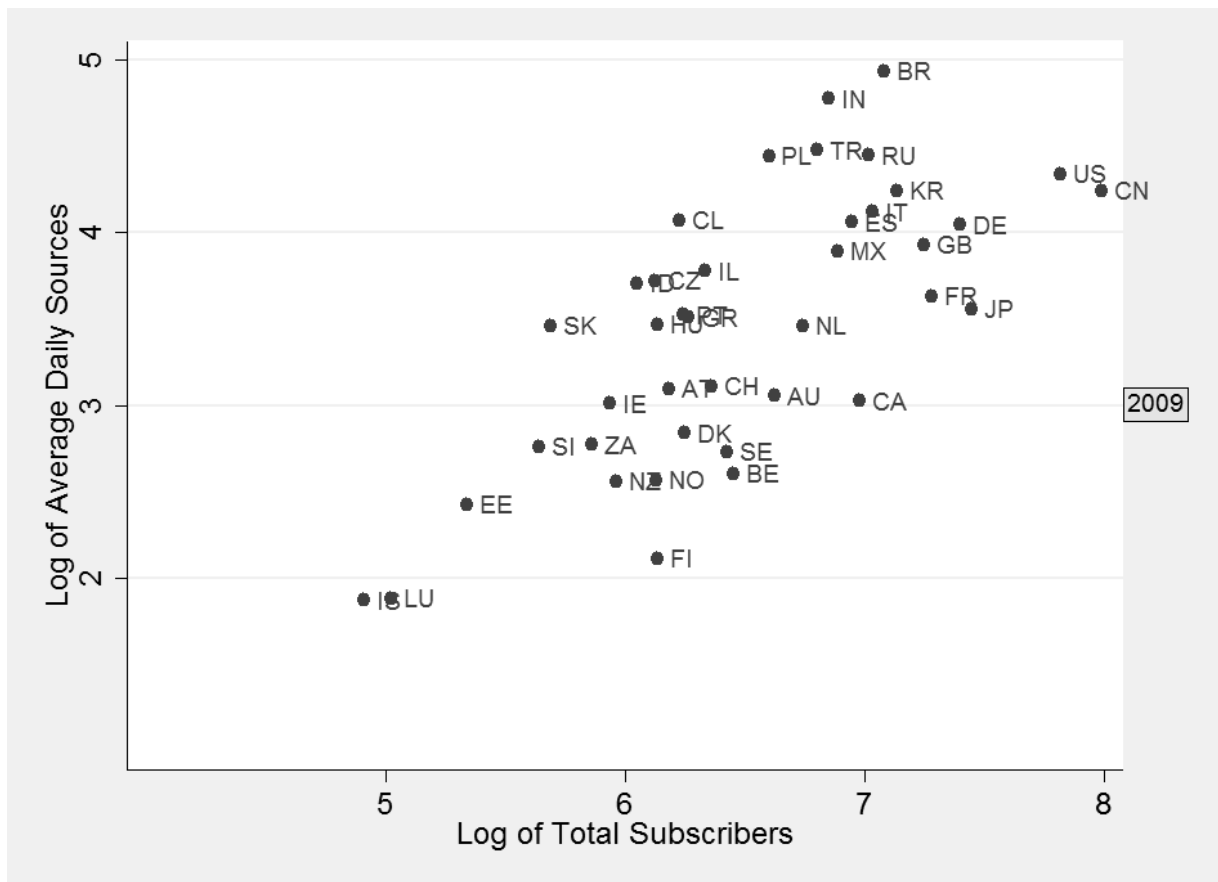
* See Appendix 2 for country codes.

Figure 11.   Average infection rates of ISPs across wider OECD area (2009)
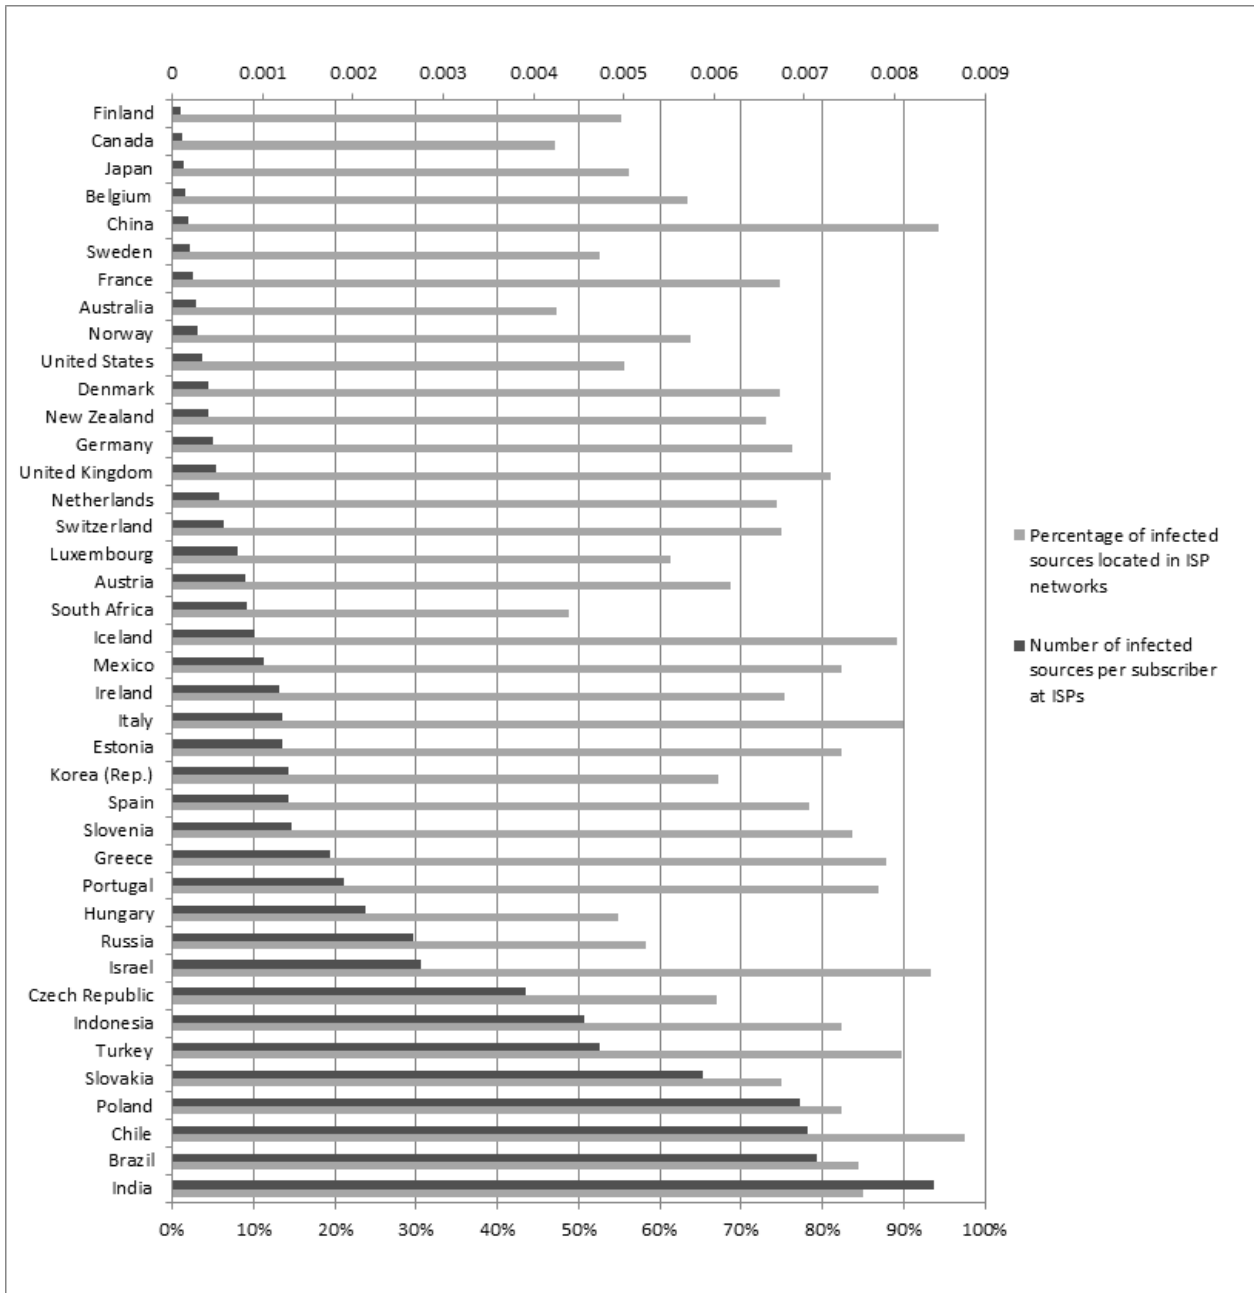
Table 4. Summary empirical results for number of infected machines per subscriber
(all entries describe the partial effects of the factor)

Factor	Main finding
Size of ISP (number of subscribers)	Large ISPs have slightly fewer infected sources per subscriber but the effect is robust and independent of model specification
Size of ISP (market share)	No relationship between market share of an ISP in its home market with the number infected sources per subscriber
Competitive pressure (average revenue per subscriber)	No relationship between ARPU of an ISP with the number of infected sources per subscriber
Status as cable modem provider	ISPs that are cable modem providers on average have fewer infected sources than DSL providers, also a robust finding
LAP membership	Membership of the London Action Plan is weakly associated with lower infection rates, an effect that is mediated by other factors such as the level of software piracy and education in a country
Cybercrime Convention membership	Countries that have joined the Cybercrime Convention have on average fewer infected sources, but the effect is statistically not significant in the regression models
Piracy rate	Software piracy has an ambiguous effect; it is sometimes associated with higher, sometimes associated with lower infection rates, varying by model specification
Education	Education shows an unambiguous effect: countries with higher education have fewer infections
Bandwidth	No relationship between average bandwidth of subscriptions in a country and the number of infected sources per subscriber for ISPs in that country

CONCLUSIONS AND POLICY IMPLICATIONS

This study set out to examine the factors that influence botnet infections using spam data as a proxy. From this knowledge we hoped to get a better understanding of options to mitigate the problem. The framework developed for the empirical analysis can also be used to identify ways to improve the performance of the overall system. Two principal options for policy intervention exist:*i)* to design measures that change the institutional and legal framework in which ISPs operate (*e.g.* measures supporting national and international law enforcement) and *ii)* to specifically target individual actors in the ICT system (*e.g.* end users, ISPs, software vendors, cybercriminals). These two options can also be combined in hybrid models, for instance public private partnerships in which both legal and corporate measures are pursued. Which particular instrument or which combination of instruments is best used depends on several criteria, including the feasibility of using it, its direct and indirect costs, and the effectiveness of the instrument. Our study does not systematically examine *all* feasible policy measures but sheds nonetheless light on some of the principal alternative approaches. Rather than drawing specific policy conclusions, which this report was not commissioned to do, this section seeks to offer insights and lessons from the empirical research that are relevant for policy makers. Before we return to this question, main findings will be summarised.

Box 11. Highlighted Findings

Throughout the report, we have used boxes to summarise our approach and the most salient findings in language that contains as few technical and statistical terms as possible. These are the main conclusions:

- The bulk of all infected machines worldwide is located in networks of well-known, legitimate ISPs in the wider OECD area – by which we mean the 33 members, plus two "accession candidates" and five OECD "enhanced-engagement" countries. Just 50 ISPs account for around half of all infected machines.

- Botnets in the wider OECD area are more or less stable, in other countries they are increasingly recruiting infected machines into the overall population of botnets

- Even good ISPs are likely tackling only a fraction of the bots

- The number of customers is the strongest factor influencing the number of infected machines in an ISP network

- ISPs of similar size and operating under similar institutional conditions can have a tenfold difference in the number of infected machines

- Attacks change rapidly, but the performance of ISPs is quite stable over time

- The most-infected ISPs are distributed across the wider OECD area and include ISPs of all sizes

- Large ISPs have, on average, fewer infected machines per customer than small ISPs

- Governmental efforts seem to help reduce infection rates at ISPs

---

The empirical data used in the study consists of spam messages that were captured by a spam trap in the period 2005-2009 – around 109 billion messages from 170 million different sources. Because around 80-90 % of all spam is issued by botnets, the origin of a spam message is very likely to indicate the presence of an infected machine. We analysed these origins with descriptive and analytical statistical

methods. To get a better understanding of the issues, we started with descriptive analyses, mapping in detail the origins of spam messages by country and ISP. This allowed an assessment of the share of infected machines located in the networks of different sets of ISPs as compared to other networks, such as hosting providers, large corporate networks, university networks, webmail providers. We found that ISPs consistently account for the bulk of all infected machines worldwide. The 200 ISPs that hold the lion's share of the access markets in the wider OECD area harbor over 60 % of all infected machines – and over 80 % of all infected machines located in those 40 countries themselves.

Combined with operational data for the ISPs, we could derive metrics, such as the average number of infected sources per user, that were then used as dependent variables to be explained by predictors related to the institutional framework, operational characteristics of ISPs, and other control variables. With regard to the control variables, we found that characteristics of the user base matter. In countries where consumers are more likely to use pirated software, we find higher botnet activity. The level of education, as a proxy for technical competence, is associated with lower levels of botnet activity. Higher average connection speeds are, however, not associated with higher levels of botnet infections, as is often presumed. In fact, we find the reverse: that high connection speeds are associated with lower botnet activity. A number of other control variables, such as the income level of a country, did not turn out to be significant factors in explaining the average number of infected sources per subscriber.

Specific policy lessons have to be derived with caution and judgment. For one, spam, malware, and botnets are dynamic phenomena. History tells us that every fortification of information security will trigger adaptations in attack strategies. Likewise, any reduction of the intensity of attacks may tempt users to reduce security investment. Both effects imply that the emergent level of security at the sytem level may respond less to policy measures than hoped. Our data point to considerable inertia in the system, which could be seen as one outcome of these effects. The findings reported in this study are based on past data and are only valid predictors of future events if the overall patterns remain in place. The five years of observations seem to indicate, however, that despite new forms of malware and new attack strategies the overall emerging patterns are fairly robust. Lastly, while the study is based on detailed data, lots of information that would be required to formulate coherent and effective policies is not systematically collected or not in the public domain. This is particularly true for information on damages from breeches of information security and for data on ISP-level security measures that would help assess which firm-level strategies are effective. The findings of the study need to be interpreted with these caveats in mind.

We found evidence to support the idea that broad governmental efforts to improve cybersecurity are associated with lower levels of botnet activity. We used membership of two international initiatives – the London Action Plan (LAP) and the Council of Europe's Convention on Cybercrime – as proxy for how active governments have been in the area of cybersecurity. The sign of the parameters was not consistent across all three modeling approaches and not all parameter estimates were statistically significant. Nonetheless, a pattern emerges that suggests that the actions undertaken by governments that have joined these international activities make a difference. This evidence must not be mistaken as causal relation in the sense that joining LAP or the Cybercrime Convention alone will reduce botnet activity. Rather, it suggests that organisation membership contributes to other government and non-government measures that have mitigating effects on botnets.

LAP membership, if taken in isolation in a bivariate analysis, was associated with lower infection rates among ISPs in those countries. Likewise, a simple regression analysis without interaction terms reveals that countries that are LAP members have lower infection rates. Qualitative evidence from selected countries supports this finding, as governements known for their active engagement of ISPs in the area of security – most notably, Japan, Australia and Finland – display better performance. Our data also confirm that the infection rates of the ISPs in those countries are lower than average. The empirical picture becomes somewhat less clear once the data is examined for possible joint effects of institutional variables.

This was done in the pooled and panel specifications by introducing interaction terms. In these model specifications, LAP membership taken by itself had a positive sign. This could be the outcome of an endogeneity problem, as selected countries with a more severe information security problem may have a higher incentive to join international co-operation efforts to mitigate the problem. However, LAP membership interacted with other variables in ways that mitigate infections. Countries that are LAP members and also have high education levels, for example, exhibit lower infection rates.

On the other hand, membership in the Cybercrime Convention always had a negative sign (*i.e.* countries that were members exhibited lower incidences of botnet infection) but the parameter estimates were generally not statistically significant. As institutional variables often affect a phenomenon jointly, several interaction terms were also tested and found to have an effect. These instititional factors cannot be seen as sufficient instruments to increase cybersecurity. We observe substantial variability among ISPs subject to the same set of institutional incentives. Thus, such measures at the institutional level, while possibly necessary conditions to enhance security, are, taken by themselves, not sufficient.

Regarding the ISPs a key finding is the degree of stability of the set of ISPs that are harboring the highest number of infections both in absolute and in relative terms. The great variability between ISPs even within one set of institutional circumstances is indirect evidence for the fact that they have a considerable degree of discretion as to how they repond to security threats. We also found that several of the factors that have been considered as important in explaining ISP security performance do not withstand empirical scrutiny. Average revenue per customer did not make a difference, which leads us to conclude that the competitive intensity of an ISP's market environment does not have a direct influence on security performance. We also tested the claim that large ISPs perform worse than smaller ones. Some experts have argued that large ISPs are less subject to peer pressure. Our data suggests that this is incorrect. In fact, large ISPs perform slightly better than average (measured by the number of infected sources and spam volume per subscriber). Market share of an ISP in its home country was not associated with worse performance either.

One speculative reason why large ISPs actually do slightly better may be that their size forces them to introduce automation in incident response and abuse management. A similar mechanism may explain why we found that cable providers did slightly better than DSL providers, especially among smaller ISPs. The management of cable networks often include automated systems and these technologies might reduce the cost of dealing with infected machines. Given the ongoing advances in technology, including botnet mitigation solutions, the difference between cable and DSL may disappear in the immediate future. Our findings do imply, however, that automation is likely to be a critical part of scaling up ISP efforts.

Our findings lend direct and indirect support to the view that ISPs are important potential control points. Not only do the legitimate, established ISPs harbor a large share of all infected machines, they also vary widely in their performance, which suggests that some have adopted more effective practices than others, even when operating under similar market and regulatory conditions. This implies that ISPs have leeway to increase their efforts, that security performance is not dictated by external conditions.

From a policy perspective, the finding that a relatively small number of ISPs is associated with a large share of total spam activity is relevant. Although these ISPs are not themselves the origin of botnet infections, they play an important role in the chain from cybercriminals to the targets of botnet attacks. The study uncovered a specific pattern that suggests that the chances of devising meaningful forms of private and public sector measures might be higher than commonly thought. On a global scale, between 4 000 and 100 000 entities can be attributed to the class of ISPs. However, we found that the distribution of infected machines is highly asymmetric. Just 50 ISPs accounted for over half of all infected sources on a global scale. Such skewed (powerlaw) distributions are familiar from many Internet-related phenomena (*e.g.* the size distribution of nodes).

To put it differently: the number of actors needed to create an impact on botnets is smaller than expected. It would be extremely difficult to bring about collective action among many thousands of ISPs located in over a hundred countries, even if ISPs were to be a more effective control point than the billion to billion-and-a-half end users. Furthermore, the most critical actors are larger, well-established corporations. It may be easier to design public policy measures and implement them for this group of ISPs, whether such measures are government interventions or forms of public-private sector co-operation. Such measures would be much more difficult if large numbers of small ISPs that are often shortlived and difficult to survey were involved. Many of these small ISPs are difficult to reach with collaborative or regulatory efforts. Even if they were interested in co-operation, the transaction cost of bringing large numbers of players into the fold may be very high.

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

While ISPs appear to have considerable discretion, their incentives are also shaped by external conditions. In retail ISP markets competition is primarily driven by price and in many countries price competition is fierce. Even if price does seem to have no significant influence on security performance, from an ISP's point of view, policy measures that affect costs (and all do directly and indirectly) are unfunded mandates and may be difficult to realise given this competitive environment. Thus, it may be necessary to think about innovative funding schemes for such programmes. Moreover, even if consumers cared about security, there are no adequate market signals that could reliably guide them towards better performing ISPs. Establishing a trusted rating system might be a tool to assist consumers in this regard. Most industry insiders lack such signals as well, except for the anecdotal evidence and speculative claims about the performance of this or that ISP that are bandied among the members of the security community. Current efforts to bring about collective action – through industry self-regulation, co-regulation, or government intervention – might initially achieve progress by focusing on the set of ISPs that together have the lion's share of the market.

The pattern we uncovered directs policy attention to established ISPs, not to the so-called grey or rogue providers which have dominated the news reports on the fight against botnets and malicious software. The reason for this is clear: rogue ISPs play a key role in the botnet economy, except for one crucial element, namely providing access to the millions of infected machines. The latter is where the legimate ISPs come in (see Box 14). Of course, there are no permanent solutions. If the 50 ISPs we identified would ramp up their efforts, the problem might migrate elsewhere. However, it is much more difficult to migrate a network of millions of infected machines than to migrate the C&C servers or other ancillary services. Furthermore, cleaning up infected machines has intrinsic value in that it protects citizens in those countries. A vulnerable machine cannot differentiate what malware it falls victim to. It exposes its owner to a range of security risks that is more diverse than becoming part of a botnet.

To conclude: the reported research is only a first step. It can be extended and deepened in several directions. First, the lack of data on important aspects of ISP performance and actions to combat botnets was a serious constraint of the study. It would be very valuable to collect such data on a systematic basis.

As some of the information is probably available in a proprietary fashion, producing it may be possible with modest effort. A public-private sector task force might be a feasible way to define the metrics that should be included. Second, the data collected for this project could be mined more systematically using shorter time periods to get a better sense of the responses of stakeholders to security incidents. In the present context, such a more finely grained perspective could not be adopted because some data utilised in the statistical analysis was only available on an annual basis. However, a more finely structured set of observations lends itself to methods of time series analysis. Third, it would be interesting to examine the causes of the differences in ISP performance – in other words, what security practices are associated with lower infection rates. To this end, original data collection will be necessary including surveys and in-depth structured interviews.

Box 12. Focus on legitimate ISPs, not just rogue ISPs

So-called 'rogue ISPs' provide many illegal services, such as distributing malware, hosting phishing sites and harboring the command and control (C&C) infrastructure of botnets, *i.e.* the servers that give the infected machines their instructions. The latter has, until now, been the key focus of the fight against botnets. When the small hosting provider McColo was forced offline, it had a massive, if short-lived, impact on the global spam volume because the provider hosted the C&C servers of a few large botnets. In addition to the McColo case, there have been other remarkable succes stories of taking down C&C infrastructure.

While it is important to pursue rogue ISPs in the fight against cybercrime, our research suggests we also need to recognise the role of legitimate ISPs. The focus on taking down C&C or other services of rogue ISPs, rather than the infected machines themselves, has serious limitations.

First of all, the successes are short-lived. It took spammers a few months to recover from the McColo shutdown, but they did recover and take spam volumes to new record heights. More recent successes were even more short-lived, as attackers innovate to make botnet C&C increasingly resilient. Many insiders worry about the moment when most botnets will use peer-to-peer technology for C&C, rather than centralised servers. There are very few ideas on how to respond to such a development.

The second reason why we should focus on established providers is simply this: whatever services rogue ISPs provide to the botnet economy, that economy would not exist without the actual bots, the millions of infected machines. In that sense, the large number of infected machines is a critical part of what criminologists call the "opportunity structure" of cybercrime.

Third, our study suggests that the millions of infected machines are a more inert part of the botnet economy, while the other parts of the problem are much more dynamic. Attackers can quickly move around C&C servers or adapt their technology to make them more resilient. This kind of agility is a lot more difficult to envision for large numbers of infected machines. We assume that this is why we found such stable ISP infection rates over time, notwithstanding all the changes in and around botnets. Such stability seems more favourable to designing effective countermeasures.

Fourth, and last, any success that might be achieved against rogue ISPs does little to clean up the actual infected machines. When a botnet is 'beheaded' by taking down the C&C servers, it leaves in place all of the infected machines. The owner of the botnet may succeed in setting up new C&C servers somewhere. Or competing attackers may recuit the infected machines into different botnets. All of this does very little to protect the owners of these machines. While the effectiveness of taking down the infected machines themselves can be debated, at least it has the benefit of directly contributing to the security of end users.

APPENDIX 1
TRIANGULATION OF OUR DATA SOURCE WITH INDUSTRY SOURCES

From the annual reports of various leading security service providers, we extracted the list of countries that sent out most spam in each year, as well as the percentage of spam volume associated with each country. We then compared the reports to each other and to our own data. The result is presented in Figures 12 for the various years – except for 2005, as the reports from that year did not include these numbers.[26] We also tested statistically how similar the sources were, by calculating the correlation among the reported numbers. These correlations are presented in Table 5.

Both approaches clearly demonstrate that our data are congruent with those provided by the security companies. The figures we found are within the range of those reported by other sources. Also, the correlation between our data with the data from other sources is high – or at least higher than the correlations among some of those sources themselves.

Table 5. Degree of similarity among data sources on spam (for 2006-2009)*

```
             | Kaspersky  Symantec Sophos   Xforce  Our data
-------------+-------------------------------------------------
   Kaspersky |  1.0000
    Symantec |  0.3946    1.0000
      Sophos |  0.3734    0.8986   1.0000
      Xforce |  0.4069    0.6620   0.7494   1.0000
    Our data |  0.4866    0.6901   0.7982   0.8417   1.0000
```

* Based on Spearman's rank correlation coefficient, a non-parametric measure of association.

Figure 12.   Countries issuing most spam 2006-2009
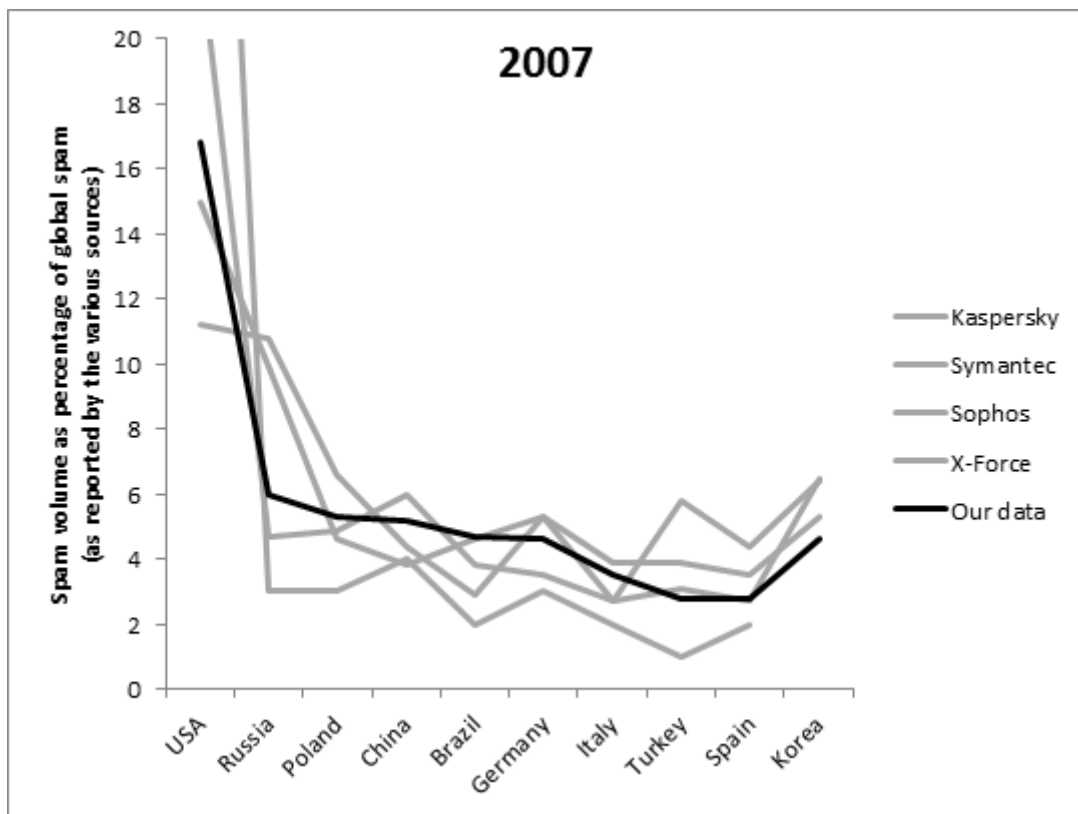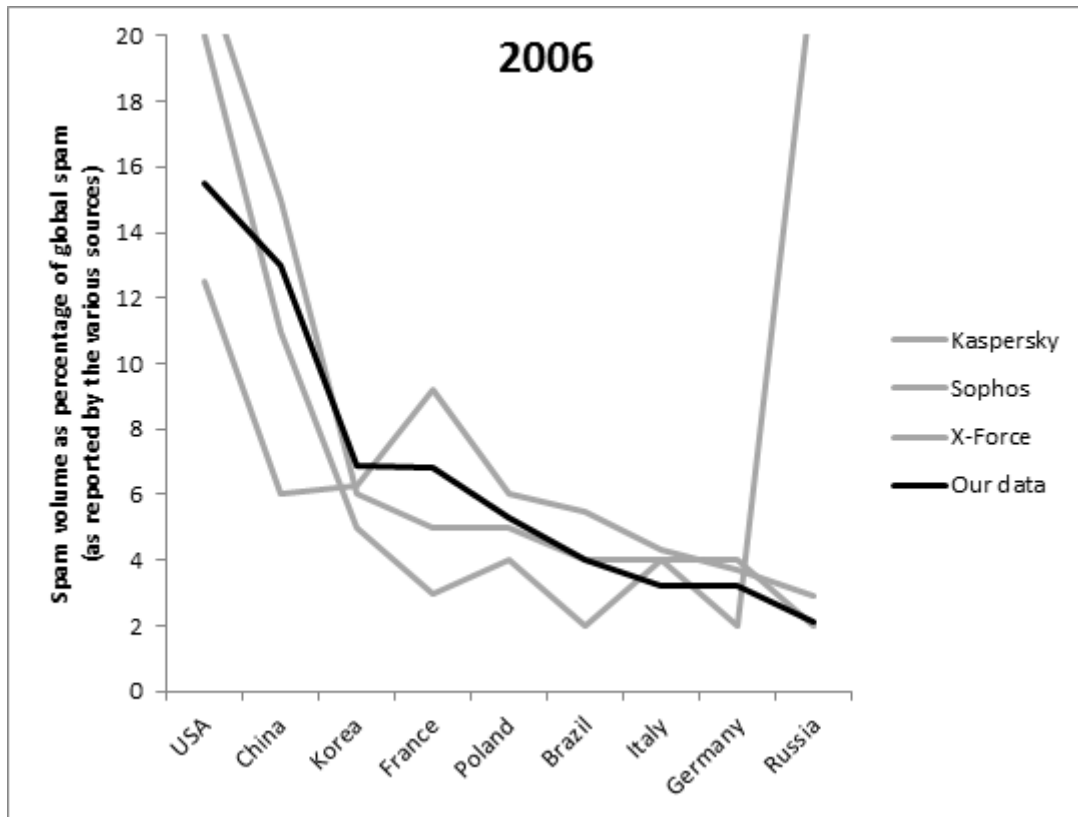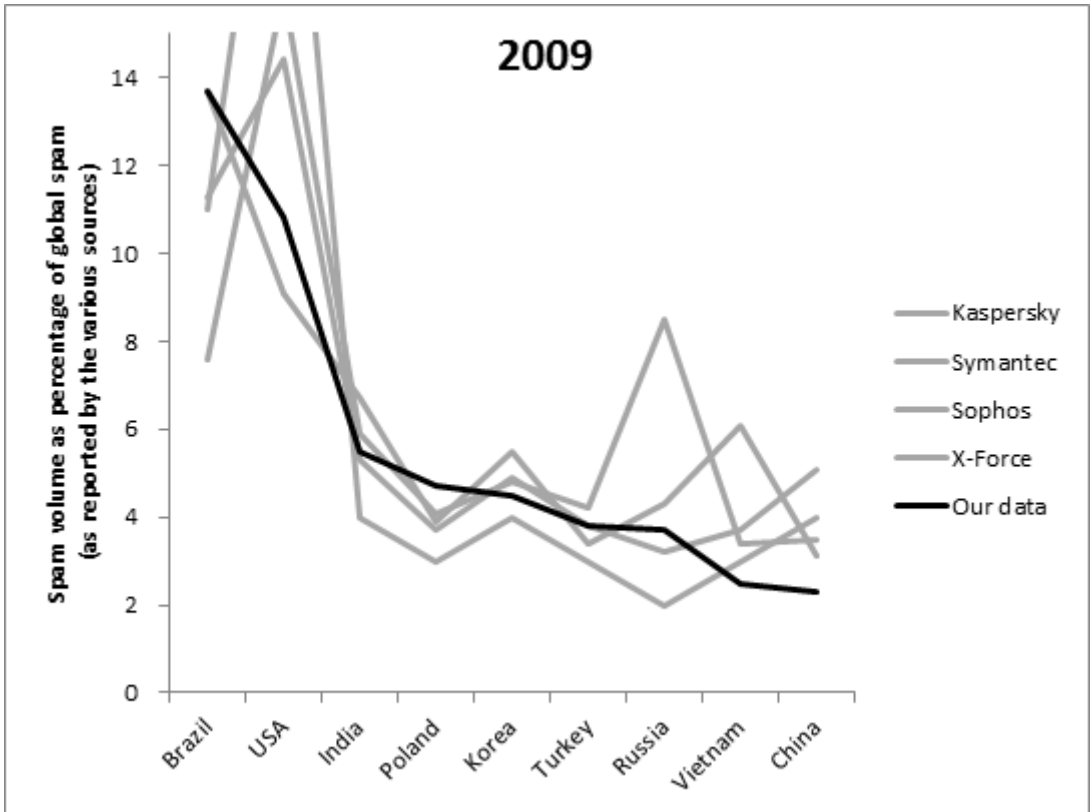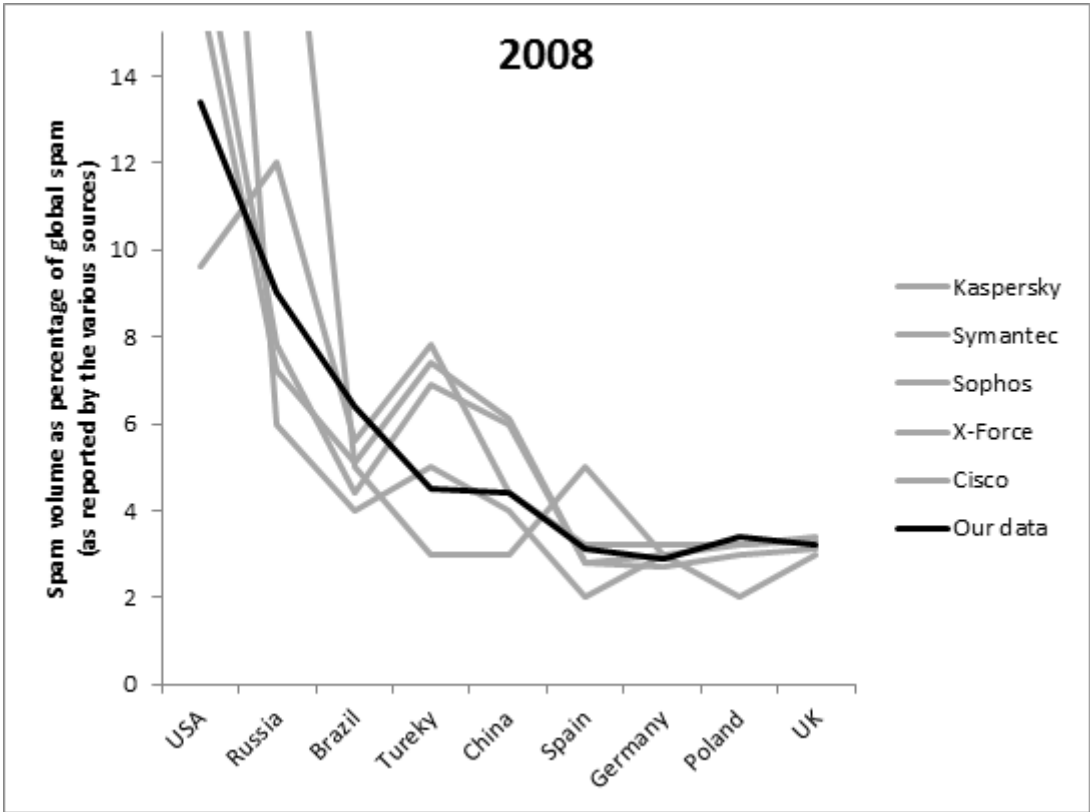
APPENDIX 2
LIST OF THE COUNTRIES AND COUNT OF ISPs INCLUDED THE FINAL DATASET

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

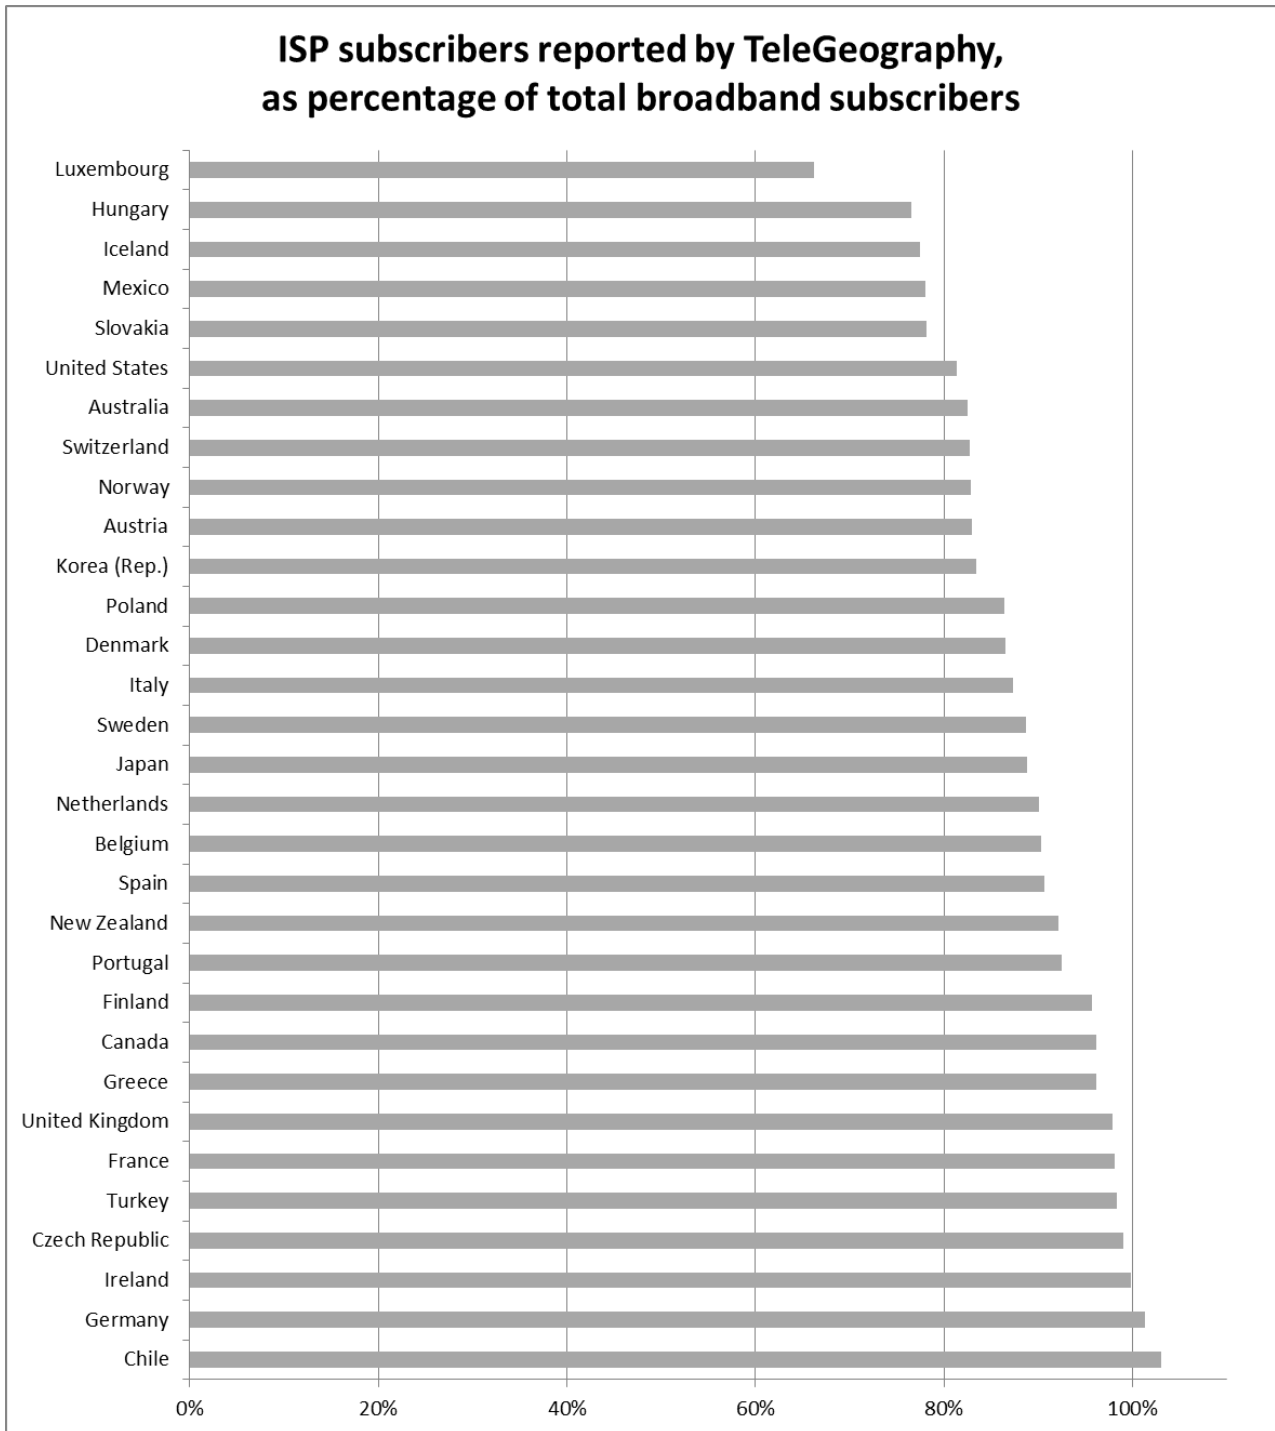**APPENDIX 3**
CROSS CHECKING THE MARKET DATA ON ISPs

To cross check the extent to which the TeleGeography database covers the broadband market in each of the countries in the wider OECD area, we have done the following: first, for each country, we add up all the subscriber numbers for the ISPs listed in the TeleGeography database. This figure gives us the total number of broadband subscribers in that country that the TeleGeography database accounts for. We then compare this figure with several publicly available sources on the total number of subscriptions in a variety of countries, namely those of the OECD[27], ITU[28], and World Bank[29]. These data sources have their own shortcomings, as they are based on self-reporting by governments. Still, we believe it provides a useful cross check on the subscriber data used in this study.

The results of this comparison using the OECD reported broadband statistics for 2009 are presented in Figure 13. The OECD statistics are widely seen as the most accurate of the publicly available sources. As can be seen from the figure, the ISPs listed by the TeleGeography database account for over 80 % of the total broadband access market in the majority of the countries – on average, 89 % of the market is covered.

The World Bank statistics for 2008, which is the most recent year currently available, and the ITU Internet Statistics for 2009, produced similar results. One notable inconsistency is regarding the subscriber totals for Israel, where the TeleGeography data accounts for more than 120 %. This is based on ITU statistics, as Israel is not reported in the OECD data. One explanation for this might be that Israeli ISPs offer satellite subscriptions to many users in other countries (*e.g.* Nigeria). These subscriptions would then show up in the operator-level statistics but not country level statistics.

A final cross check was performed directly with Internet service providers in the Netherlands. These ISPs provided market figures directly, which were compared with the TeleGeography figures. The congruency was very high with the difference being approximately 1 %.

Figure 13.   Market share of ISPs included in the study

APPENDIX 4
COMPENSATING FOR KNOWN LIMITATIONS IN INTERNET MEASUREMENTS

Our approach allows us to robustly establish the number of infected machines in ISP networks. It has certain limitations, however, that need to be compensated. The effects of three technical issues need to be taken into account when interpreting the data: the use of Network Address Translation (NAT), the use of dynamic IP addresses with short lease times, and the use of port 25 blocking. The key issue is to understand how these technical practices affect the number of machines that are represented by a single unique IP address.

NAT means sharing a single IP address among a number of machines. Home broadband routers often use NAT, as do certain other networks. This potentially under-represents the number of infected machines, as they all show up as a single address. Dynamic IP addresses with short lease times implies that a single machine will be assigned multiple IP addresses over time. This means a single infected machine can show up under multiple IP addresses. As such, it over-represents the number of infected machines. Both these practices counteract each other, to some extent. This limits the bias each of them introduces in the data, but this does not happen in a consistent way across different networks.

This is a classic problem in the field of Internet measurement: how many machines are represented by a single IP address? Ideally, one IP address would indicate one machine. But reality is more complicated. Over an extended time period, a single address sometimes indicates less than one machine, sometimes more than one. This varies across ISPs and countries. Earlier research by Stone-Gross *et al*. (2009) has demonstrated that in different countries, there are different ratios of unique IP addresses to infected machines – the so-called "churn rates".

We have two ways to robustly control for the potential bias that these churn rates introduce in our data. First, we look at the volume of spam in addition to the number of unique sources. If there are several machines behind a single IP address, the spam volume is also several times higher that that of a single machine. If there is one machine behind many IP addresses, the spam volume is proportionally lower for each address. We have calculated the ratio of unique sources to spam volume in our data. The Spearman correlation between the churn rates reported by Stone-Gross *et al*. (2009) and ratios we calculated is very high, namely 0.88. This resemblance suggests that spam volume can indeed control for churn.

A second way to control for it is to use shorter time scales when counting the number of unique IP addresses in a network. The potential impact of churn is very limited on shorter time scales,. Research by Moore *et al*. (2002) found that churn starts to affect the accuracy of IP addresses as a proxy for machines on timescales longer than 24 hours. We therefore worked with a timescale of 24 hours. All our analyses are based on the daily average number of IP addresses sending spam from an ISP network.

A final limitation is the use of port 25 blocking by ISPs. The effect of port blocking is that infected machines can no longer directly send e-mail to the wider Internet, but have to go through the ISP's outgoing e-mail servers. This affects both the number of sources as well as the spam volume. The ISP's network may harbor thousands of infected machines, but they can no longer reach the spam trap directly and thus do not reveal their IP address through spam distribution. There is one important way in which the attackers themselves compensate for this problem: when the bots notice they cannot connect anymore via port 25, they start to redirect their spam through the ISP's official outgoing e-mail servers. In various cases

where port blocking was introduced, we saw that it led to a brief reduction of outgoing spam, only to return to the previous spam volume within about a month. It is difficult for the ISP to prevent this from happening, as each bot sends out a relatively low level of spam, and thus rate limits and similar controls often do not pick up on it. This adaptation of the spam bots allows us to use spam volume to cross check our findings. It is not perfect, however. Port blocking is an unavoidable limitation to our data. If the spam volume remains consistently lower, port blocking obscures the presence of infected machines. That being said, the effect of the bias is not wholly unreasonable. The ISPs that adopt port blocking improve their ranking in terms of botnet activity compared to those that don't – which is not without merit, given that the measure of port blocking is part of many guidelines on best security practices for ISPs and that it cuts into the criminal business model of spammers.

For all the analyses we discuss in this paper, we have always checked whether the pattern we found persisted across all three variables: daily average number of unique IP addresses sending spam in a network, total number of unique IP addresses sending spam in a network per year, and the volume of spam from a network per year. That way, we can compensate for the various measurement issues. Patterns that hold across these different measurements can be said to be robust and valid. For the sake of brevity, we focus our discussion on the average daily number of unique sources. When spam volume or the total number of unique sources per year show a different pattern, we explicitly include it in the discussion. Where they are not mentioned, they are consistent with the findings as reported here.

APPENDIX 5
DATA AND DATA SOURCES

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

APPENDIX 6
DESCRIPTIVE STATISTICS

```
    Variable |     Obs        Mean    Std. Dev.        Min        Max
-------------+--------------------------------------------------------
    unq_srcs |     932    189376.1     509729.7         13    6420173
  src_persub |     932    .1843209     .2046522      .0001     1.1329
    avg_uips |     932    1624.425     4422.309          0      67717
 auip_persub |     932    .0018221     .0024104          0      .0178
   spam_msgs |     932    6.21e+07     1.62e+08       7679    2.50e+09
-------------+--------------------------------------------------------
 spam_persub |     932    76.75756     132.4606       .027   1903.709
   total_sub |     932     1439564      3757275       3000    5.35e+07
market_share |     699    .1811987     .1995256      .0005     1.2358
  rev_persub |     148    4471.662     5159.659   182.1285   42768.34
   srv_cable |     828    .3671498     .4823192          0          1
-------------+--------------------------------------------------------
     lap_mem |     932     .554721     .4972634          0          1
   cyber_mem |     932    .7113734     .4533672          0          1
 piracy_rate |     930    40.32688       16.864         20         87
     educ_ix |     932     .947133     .0616795       .632       .993
     int_bpp |     383    14327.88     17031.71   190.8559   92832.46
```

60

APPENDIX 7
PAIR WISE CORRELATIONS BETWEEN THE INDEPENDENT VARIABLES

```
              | total_~b market~e rev_pe~b srv_ca~e  lap_mem cyber_~m piracy~e
--------------+-------------------------------------------------------------
   total_sub |  1.0000
market_sha~e |  0.2516   1.0000
  rev_persub | -0.0784   0.2254   1.0000
   srv_cable | -0.1077  -0.2009  -0.1953   1.0000
     lap_mem |  0.1424  -0.1769  -0.0491   0.0639   1.0000
   cyber_mem | -0.0637  -0.0552   0.1067   0.0494   0.1249   1.0000
 piracy_rate |  0.0657   0.1085   0.0620  -0.0604  -0.3829  -0.5767   1.0000
     educ_ix | -0.0931  -0.0765  -0.3094   0.1313   0.3157   0.4515  -0.5788
     int_bpp | -0.0285  -0.0416  -0.1190  -0.0141   0.2531   0.4692  -0.5171

              |  educ_ix  int_bpp
--------------+------------------
     educ_ix |  1.0000
     int_bpp |  0.3203   1.0000
```

61

REFERENCES

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

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

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

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

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

Clayton, R. (2010), *Might Governments Clean-up Malware?* Ninth Workshop on the Economics of Information Security (WEIS 2010) Harvard University. Available online at http://weis2010.econinfosec.org/papers/session4/weis2010_clayton.pdf.

Clover, C. (2009), *Kremlin-backed group behind Estonia cyber blitz*. The *Financial Times* website. Available online at www.ft.com/cms/s/0/57536d5a-0ddc-11de-8ea3-0000779fd2ac.html.

Federal Trade Commission (2007), *Spam Summit: The Next Generation of Threats and Solutions*. Available online at www.ftc.gov/os/2007/12/071220spamsummitreport.pdf.

Fox, J. (2007). *Consumer Reports: Putting Consumers Back in Control*. Available online at www.ftc.gov/bcp/workshops/spamsummit/presentations/Consumers.pdf, November 25, 2007.

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

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

Krebs, B. (2008), *Spam Volumes Drop by Two-Thirds After Firm Goes Offline*. *Washington Post* Security Fix weblog. Available online at http://voices.washingtonpost.com/securityfix/2008/11/spam_volumes_drop_by_23_after.html.

Messagelabs (2007), *MessageLabs Intelligence: 2006 Annual Security Report*. Available online at www.messagelabs.com/mlireport/2006_annual_security_report_5.pdf.

Messagelabs (2009), *MessageLabs Intelligence: 2008 Annual Security Report*. Available online at www.messagelabs.com/mlireport/MLIReport_Annual_2008_FINAL.pdf.

Messagelabs (2010), *MessageLabs Intelligence: 2009 Annual Security Report*. Available online at www.messagelabs.com/mlireport/2009MLIAnnualReport_Final_PrintResolution.pdf.

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

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

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

OECD (2010), *The Economic and Social Role of Internet Intermediaries*.Paris, Organisation for Economic Co-operation and Development. Available online at www.oecd.org/dataoecd/49/4/44949023.pdf.

Perrow, C. (2007), *The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters*. Princeton, NJ, Princeton University Press.

Peterson, P. (2007), *Evolving Methods for Sending Spam and Malware*. Federal Trade Commission Spam Summit. Available online at www.ftc.gov/bcp/workshops/spamsummit/presentations/Evolving-Methods.pdf.

Sandvine (2004), *Trend analysis: Spam trojans and their impact on broadband service providers*. Available online at www.sandvine.com/general/getfile.asp?FILEID=13.

Soper, M. E. (2009), *Conficker worm shuts down French and UK Air Forces*. Maximum PC. Available online at www.maximumpc.com/article/news/conficker_worm_shuts_down_french_and_uk_air_forces

Sophos (2010), *Security Threat Report: 2010*. Available online at www.sophos.com/sophos/docs/eng/papers/sophos-security-threat-report-jan-2010-wpna.pdf.

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

Spring, T. (2005), ″Spam Slayer: Slaying Spam-Spewing Zombie PCs″. *PC World*. Available online at www.pcworld.com/article/121381/spam_slayer_slaying_spamspewing_zombie_pcs.html.

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

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

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

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

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

NOTES

1. OECD member countries are Australia, Austria, Belgium, Canada, Chile, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Israel, Italy, Japan, Korea, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, United Kingdom and United States. Estonia and the Russian Federation are candidate to accession. Estonia has been invited to become a member in May 2010. Brazil, China, India, Indonesia and South Africa are Enhanced Engagement countries.

2. The authors of this study are Michel van Eeten (m.j.g.vaneeten@tudelft.nl), Johannes M. Bauer (bauerj@msu.edu), Hadi Asghari (h.asghari@tudelft.nl), Shirin Tabatabaie (s.tabatabaie@tudelft.nl). M. van Eeten, H. Asghari and S. Tabatabaie are affiliated with the Delft University of Technology, Faculty of Technology, Policy and Management, the Netherlands. J. Bauer is affiliated with Michigan State University, Department of Telecommunication, Information Studies, and Media, United States.

   The authors gratefully acknowledge the generous support provided by Dave Rand of TrendMicro who provided the raw data from a spam trap which form the basis for the study. The report also benefitted greatly from feedback provided by anonymous reviewers of an earlier draft which was presented at the Workshop on the Economics of Information Security (Harvard University, June 7-8, 2010).

3. A recent report by the OECD defines Internet intermediaries as follows: "Internet intermediaries bring together or facilitate transactions between third parties on the Internet. They give access to, host, transmit and index content, products and services originated by third parties on the Internet or provide Internet-based services to third parties." See: OECD (2010). *The Economic and Social Role of Internet Intermediaries*. OECD. Available online at www.oecd.org/dataoecd/49/4/44949023.pdf..

4. Messagelabs (2010), *MessageLabs Intelligence: 2009 Annual Security Report*. Available online at www.messagelabs.com/mlireport/2009MLIAnnualReport_Final_PrintResolution.pdf.

5. Messagelabs (2009), *MessageLabs Intelligence: 2008 Annual Security Report*. Available online at www.messagelabs.com/mlireport/MLIReport_Annual_2008_FINAL.pdf.

6. Messagelabs (2007), *MessageLabs Intelligence: 2006 Annual Security Report*. Available online at www.messagelabs.com/mlireport/2006_annual_security_report_5.pdf.

7. Spring, T. (2005), ″Spam Slayer: Slaying Spam-Spewing Zombie PCs″. *PC World.* Available online at www.pcworld.com/article/121381/spam_slayer_slaying_spamspewing_zombie_pcs.html.

8. Peterson, P. (2007), *Evolving Methods for Sending Spam and Malware*. Federal Trade Commission Spam Summit. Available online at www.ftc.gov/bcp/workshops/spamsummit/presentations/Evolving-Methods.pdf.

9. Federal Trade Commission (2007), *Spam Summit: The Next Generation of Threats and Solutions*. Available online at www.ftc.gov/os/2007/12/071220spamsummitreport.pdf.

10. Sandvine (2004), *Trend analysis: Spam trojans and their impact on broadband service providers*. Available online at www.sandvine.com/general/getfile.asp?FILEID=13.

11.   Ironport reported a drop in spam volume by half, Sophos by three-quarters. See Krebs, B. (2008). ″Spam Volumes Drop by Two-Thirds After Firm Goes Offline″. *Washington Post* Security Fix weblog. Available online at http://voices.washingtonpost.com/securityfix/2008/11/spam_volumes_drop_by_23_after.html. Sophos (2010). *Security Threat Report: 2010*. Available online at www.sophos.com/sophos/docs/eng/papers/sophos-security-threat-report-jan-2010-wpna.pdf.

12.   Spring, T. (2005). ″Spam Slayer: Slaying Spam-Spewing Zombie PCs″. *PC World*. Available online at www.pcworld.com/article/121381/spam_slayer_slaying_spamspewing_zombie_pcs.html.

13.   Sophos (2010). *Security Threat Report: 2010*. Available online at www.sophos.com/sophos/docs/eng/papers/sophos-security-threat-report-jan-2010-wpna.pdf.

14.   Spamhaus defines snowshoe spam as follows: "Like a snowshoe spreads the load of a traveler across a wide area of snow, some spammers use many frequently-changing IP addresses and domains to spread out the spam load in order to dilute recipient reputation metrics and evade filters." See: www.spamhaus.org/faq/answers.lasso?section=Glossary#233

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

16.   Anderson *et al.* (2008), suggest that there are 40 000 ISPs in the EU alone.

17.   For larger countries, such as the United States, we used a lower threshold of 0.1 % of total spam volume for that country. We mapped the ASNs in each country, over 900 in total, by going down the list of top spam-sending ASNs, ranked by volume, until one of the following conditions was met: *i)* 95 % of the spam originating from that country had been covered; or *ii)* the number of ASNs covered is five times the number of ISPs in that country, as listed in TeleGeography's GlobalComms database; or *iii)* the next ASN contributes less than 1 % of spam originating from that country and less than 0.01 % of spam worldwide.

18.   Moore *et al.* (2009), mention several notorious examples of rogue providers involved in the botnet economy. All of these are hosting providers, not access providers harboring the actual infected machines. See: Moore, T., R. Clayton and R. Anderson (2009). ″The Economics of Online Crime″. *Journal of Economic Perspectives* 23(3): 3-20.

19.   The coefficient of variation – basically the ratio of the standard deviation to the mean – is well above 1, for the daily average number of unique sources per subscriber. This supports the finding that there are substantial differences within the whole set of ISPs, *i.e.* it is not a matter of a few outliers.

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

21.   These statistics are published annually by the Business Software Alliance and can be accessed via www.bsa.org/country/Research%20and%20Statistics.aspx.

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

23 We used a logarithmic transformation of subscriber count because the order of magnitude of the number of subscribers is more important than the absolute number – *i.e.* we would expect security practices to differ between an ISP with 50 000 subscribers and one with 500 000, but not between ISPs with 5 million and 5.5 million subscribers.

24. With only five years of observations, though, panel data estimation has limitations as the time-series component is relatively short. Of the three methods, panel data estimation therefore is the most challenging approach for finding a empirical model that yields statistically significant parameter estimates.

25. An analysis of the error terms indicated the presence of heteroskedasticity. This means that the error terms are drawn from different distributions for different values of the independent variables. In other words, the error terms vary with one or more of the independent variable(s). Although heteroskedasticity weakens the findings it does not invalidate them. It is a possible indication that other factors that are not yet included in the model may be at work. For the pooled model, we corrected the bias in the estimated standard error by running a so-called robust regression. For the panel model, we used feasible general least squares (FGLS) and therefore took the heteroskedasticity into account but allowed for an uncorrelated error structure. Both ways (correcting standard errors and relying on FGLS) are common methods of dealing with heteroskedasticity. Due to these problems, R2 is not defined. Model specifications that do not control for these potential issues (and therefore yield less reliable parameter estimates and significance levels) suggest that slightly more than 20 % of the variance is explained.

26. Sources for the 2009 figures include: IBM Security Solutions X-Force 2009 Trend and Risk Report (released April 2010); Symantec Global Internet Security Threat - Report Trends for 2009 (published April 2010); Sophos Security Threat Report: 2010; Kaspersky Security Bulletin: Spam Evolution 2009 (available on securelist.com). For the previous years, older versions of the same reports have been used. Apart from the mentioned vendors, the security reports from Cisco, McAfee, Microsoft, PandaLabs and Trend Micro were also checked. These reports didn't contain comparable figures (*i.e.* no list of top spam-senders, or the lists weren't on an annual basis).

27. Available online at www.oecd.org/sti/ict/broadband.

28. Available online at www.itu.int/ITU-D/ICTEYE/Indicators/Indicators.aspx.

29. Available online at http://data.worldbank.org/indicator/IT.NET.BBND.P2.