Unclassified **DSTI/DOC(2008)1**

Organisation de Coopération et de Développement Économiques
Organisation for Economic Co-operation and Development

29-May-2008

English - Or. English

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY

DSTI/DOC(2008)1
Unclassified

ECONOMICS OF MALWARE: SECURITY DECISIONS, INCENTIVES AND EXTERNALITIES

STI WORKING PAPER 2008/1
Information and Communication Technologies

Michel J.G. van Eeten and Johannes M. Bauer

JT03246705

English - Or. English

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.

---

http://www.oecd.org/sti/working-papers

ECONOMICS OF MALWARE: SECURITY DECISIONS, INCENTIVES AND EXTERNALITIES

Michel J.G. van Eeten[1] and Johannes M. Bauer[2]
with contributions from Mark de Bruijne, Tithi Chattopadhyay, Wolter Lemstra, John Groenewegen, and Yuehua Wu

ABSTRACT

Malicious software, or malware for short, has become a critical security threat to all who rely on the Internet for their daily business, whether they are large organisations or home users. While originating in criminal behaviour, the magnitude and impact of the malware threat are also influenced by the decisions and behaviour of legitimate market players such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. This working paper reports on qualitative empirical research into the incentives of market players when dealing with malware. The results indicate a number of market-based incentive mechanisms that contribute to enhanced security but also other instances in which decentralised actions may lead to sub-optimal outcomes - *i.e.* where significant externalities emerge.

1 . Faculty of Technology, Policy and Management, Delft University of Technology - m.j.g.vaneeten@tudelft.nl.

2 . Quello Center for Telecommunication Management & Law, Michigan State University - bauerj@msu.edu.

ECONOMIE DU "MALWARE": DECISIONS DE SECURITE, INCITATIONS ET EXTERNALITES

Michel J.G. van Eeten[3] et Johannes M. Bauer[4]
avec les contributions de Mark de Bruijne, Tithi Chattopadhyay, Wolter Lemstra, John Groenewegen, et Yuehua Wu.

RÉSUMÉ

Les logiciels malveillants, ou "malware", sont devenus une menace sérieuse pour tout ceux dont les activités quotidiennes reposent sur l'utilisation d'Internet, qu'il s'agisse de grandes organisations ou de particuliers. Bien qu'elle trouve sa source dans un comportement criminel, l'étendue et les conséquences de cette menace sont également influencées par les décisions et les comportements d'acteurs légitimes du marché tels que les fournisseurs d'accès Internet, vendeurs de logiciels, entreprises de commerce électronique, fabricants de matériel informatique et registres, sans oublier les utilisateurs finals. Ce document reflète le contenu d'une recherche qualitative et empirique concernant les incitations des acteurs du marché lorsqu'ils sont confrontés au malware. Les résultats indiquent qu'il existe des incitations fondées sur le marché qui contribuent à augmenter la sécurité mais également des cas dans lesquels des actions décentralisées peuvent conduire à des résultats sous-optimaux, i.e. où des externalités significatives émergent.

3 . Faculty of Technology, Policy and Management, Delft University of Technology - m.j.g.vaneeten@tudelft.nl.

4 . Quello Center for Telecommunication Management & Law, Michigan State University - bauerj@msu.edu.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Malicious software, or malware for short, has become a critical security threat to all who rely on the Internet for their daily business, whether they are large organisations or home users. While initially a nuisance more than a threat, viruses, worms and the many other variants of malware have developed into a sophisticated set of tools for criminal activity. Computers around the world, some estimate as many as one in five, are infected with malware, often unknown to the owner of the machine. Many of these infected machines are connected through so-called botnets: networks of computers that operate collectively to provide a platform for criminal purposes. These activities include, but are not limited to, the distribution of spam (the bulk of spam now originates from botnets), hosting fake websites designed to trick visitors into revealing confidential information, attacking and bringing down websites, enabling so-called 'click fraud,' among many other forms of often profit-driven criminal uses. There are also reports that indicate terrorist uses of malware and botnets. This report, however, focuses primarily on malware as an economic threat.

While originating in criminal behaviour, the magnitude and impact of the malware threat is also influenced by the decisions and behaviour of legitimate market players such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. All of these market players are confronted with malware, but in very different ways. Most importantly, they face different costs and benefits when deciding how to respond to malware. In other words, they operate under different incentives.

As security comes at a cost, tolerating some level of insecurity is economically rational. From an economic perspective, the key question is whether the costs and benefits perceived by market players are aligned with social costs and benefits of an activity. In certain situations, the security decisions of a market player regarding malware may be rational for that player, given the costs and benefits it perceives, but its course of action may impose costs on other market players or on society at large. These costs are typically not taken into account by the market player making the initial decision, causing an "externality." Externalities are forms of market failure that lead to sub-optimal outcomes if left unaddressed. In the presence of externalities, Internet-based services may be less secure than is socially desirable. This study has primarily an empirical and analytical focus and intends to document these effects. Whereas new policies may be required to address these problems, developing recommendations for such policies is outside the scope of this report.

We set out to identify externalities by analysing the incentives under which a variety of market players operate when dealing with malware. The core of the report is made up of a detailed discussion of the outcomes of a qualitative empirical field study. In the course of 2007, we conducted 41 in-depth interviews with 57 professionals from organisations participating in networked computer environments that are confronted with malware. Based on this unique data, we identified the key incentives of ISPs, e-commerce companies (with a focus on financial service providers), software vendors, registrars and end users.

The results indicate a number of market-based incentive mechanisms that contribute to enhanced security but also other instances in which decentralized actions may lead to sub-optimal outcomes – *i.e.* where significant externalities emerge. A pressing question is whether the response to malware of actors in information and communication markets is adequate or whether improvements are possible. Pointing to a variety of reports that show increases in malicious attack trends, one might conclude that markets are not responding adequately. Our analysis revealed a more nuanced picture.

With regard to the interrelationships within the information and communications-related activities, it seems that the incentives of many of the market players are reasonably aligned with minimising the effects of externalities on the sector as a whole. The incentives typically have the correct

directionality, but in a variety of cases they are too weak to prevent significant externalities from emerging. It is important to note, however, that all market players we studied experience at least some consequences of their security tradeoffs on others. There are feedback loops, such as reputation effects, that bring some of the costs imposed on others back to the agent that caused them – even if in some cases, the force of the feedback loop has so far been too weak or too localised to move behaviour swiftly towards more efficient social outcomes.

Across the value net of the different market players, three relevant situations emerge:

i) No externalities. This concerns instances in which a market player, be it an individual user or an organisation, correctly assesses security risks, bears all the costs of protecting against security threats (including those associated with these risks) and adopts appropriate counter measures. Private and social costs and benefits of security decisions are aligned. There may still be significant damage caused by malware, but this damage is borne by the market player itself. This situation would be economically efficient but, due to the high degree of interdependency in the Internet, it is relatively rare.

ii) Externalities that are borne by agents in the value net that can manage them. This concerns instances in which a market player assesses the security risks based on the available information but, due to the existence of (positive or negative) externalities, the resulting decision deviates from the social optimum. Such deviations may be based on lack of incentives to take costs imposed on others into account, but it can also result from a lack of skills to cope with security risks, or financial constraints faced by an individual or organisation. As long as somebody in the value net internalises these costs and this agent is in a position to influence these costs – *i.e.* it can influence the security tradeoffs of the agents generating the externality – then the security level achieved by the whole value net will deviate less from a social optimum than without such internalisation. This scenario depicts a relatively frequent case and numerous examples were found that confirm externalities were being internalised by other market players.

For example, the incentives of financial service providers are such that in many cases they compensate customers for the damage they suffer from online fraud. In that sense, they internalise the externalities of sub-optimal security investments of their customers as well as the software vendors whose software is exploited to execute the attacks. Many financial service providers claim they compensate all malware-related losses. If that claim is accurate, then the security level achieved by the whole value net may not be too far from the social optimum. The financial institutions bear the externalities, but they are also in a position to mitigate the size of these externalities, *i.e.* they can manage the risk through the security measures around online financial services. Within their incentive structure, it currently is more efficient to keep malware-related losses at acceptable levels, rather than to aggressively seek to reduce them. A dominant incentive is the benefits of a growing online transaction volume. Any security measure that might reduce the ease of use of online financial services may impede this growth, which implies costs that are likely to be much higher than the current direct damage from malware-related fraud (see Chapter III for more details).

iii) Externalities that are borne by agents who cannot manage them or by society at large. An individual unit may correctly assess the security risks given its perceived incentives but, due to the existence of externalities, this decision deviates from the social optimum. Alternatively, an individual unit may not fully understand the externalities it generates for other actors. Unlike in scenario two, no other agents in the information and communication value net absorb the cost or, if they do, they are not in a position to influence these costs – *i.e.* influence the security tradeoffs of the agents generating the externality. Hence, costs are generated for the whole sector and society at large. These are the costs of illegal activity or crime associated with malware, the costs of restitution of crime victims, the costs of e-commerce companies buying security services to fight off botnet attacks, the cost of law enforcement associated with these activities, and so forth. Furthermore, they may take on the more indirect form of

slower growth of e-commerce and other activities. Slower growth may entail a significant opportunity cost for society at large if the delayed activities would have contributed to economic efficiency gains and accelerated growth. A comprehensive assessment of these additional costs will demand a concerted effort but will be necessary to determine the optimal level of action to fight malware.

The most poignant cases in this category are the externalities caused by lax security practices of end users – not limited to home users, but across the spectrum up to and including large organisations such as retailers or governmental institutions. Some of these externalities are internalised by other market players that can mitigate them, most notably ISPs that can quarantine infected end users, but only to a limited extent. ISPs have incentives to deal with these problems only in so far as they themselves suffer consequences from the end user security failures, *e.g.* by facing the threat that a significant part of their network gets blacklisted. Estimates mentioned in the interviews suggest that the abuse notifications that ISPs receive concern only a fraction of the overall number of infected machines in their network.

Consequently, many externalities emanating from end user behaviour are borne by the sector as a whole and society at large. These externalities are typically explained by the absence of incentives for end users to secure their machines. It would be more precise, however, to argue that the end users do not perceive any incentives to secure their machines. While malware writers have purposefully chosen to minimize their impact on the infected host and to often direct their attacks at other targets, there is also a plethora of malware which does in fact attack the infected host – most notably to scour any personal information that can be used for financial gain. In that sense, end users do have a strong incentive to secure their machines. Unsecured machines cannot differentiate between malware that does or does not affect the owner of the machine. If the machine is not sufficiently secured, then one has to assume that all forms of malware can be present. The fact that this incentive is not perceived by the end user is an issue of incomplete information rather than a lack of incentives.

Although the research reported in this report was not designed to develop specific policy recommendations, some general concluding remarks are offered. We found many feedback loops which mitigate the externalities arising from security-reducing behaviour. All market players we studied experience such feedback, which potentially better aligns their decisions with the social optimum. We also noted, however, that in many cases these feedback loops are too weak or localised to effectively change the security tradeoffs from which the externalities emerge. In terms of policy development, a key strategy would be to strengthen the existing feedback loops and create new ones where possible. That would also keep public policy out of the realm of having to decide how secure is secure enough when it comes to defending against malware.

ACKNOWLEDGEMENTS

A study such as ours incurs considerable debt along the way. First and foremost, we thank our interviewees, who gave generously of their time. They also provided valuable comments on a draft version of this report and checked and approved the use of their quotes where appropriate. Their input is greatly appreciated. To maintain confidentiality, none of those interviewed is named in the text.

Special thanks go to our colleagues Mark de Bruijne, Wolter Lemstra and John Groenewegen in Delft and Tithi Chattopadhyay, Yuehua Wu in East Lansing. They have provided invaluable contributions in the course of this project and we have greatly benefited from the exchanges of ideas with them.

We also would like to thank Anne Carblanc, Audrey Plonk and Sam Paltridge at the OECD and Ronald van der Luit and Edgar de Lange at the Netherlands Ministry of Economic Affairs for supporting this research and for their engaging questions and comments. Selected findings from this report are included in the OECD's report on *Malicious Software (Malware)*: *A Security Threat to the Internet Economy*, developed in collaboration with the APEC Telecommunications Working Group.

We have given presentations to conferences on our findings, including the 35th Telecommunications Policy Research Conference (Alexandria, VA, September 28-30, 2007), the LAP/CNSA/MAAWG Workshop (Arlington, October 9-11, 2007), the 2007 GOVCERT conference (Noordwijk, October 18-19, 2007). Some of the very best feedback has been from the presentation of our interim findings at the meetings of the OECD WPISP and the workshops with policy makers at the Dutch Ministry of Economic Affairs.

I. INTRODUCTION

The past five years have witnessed the emergence of comprehensive efforts to improve the security of information systems and networks. A recent survey by the OECD (2005) demonstrates that governments have developed national policy frameworks as well as partnerships with the private sector and civil society to combat cybercrime. Measures include Computer Security Incident Response Teams (CSIRTs), raising awareness, information sharing, and fostering of education.

During the same period, security threats have increasingly captivated the public's attention – fuelled by new attack trends on the Internet, terrorism warnings, rising cybercrime and our growing reliance on the Internet and other communication networks in virtually all aspects of our lives. An increasingly powerful threat is posed by so-called "malware" – commonly defined as malicious software that is inserted into an information system, usually covertly, with the intent of compromising the confidentiality, integrity, or availability of the victim's data, applications, or operating system or otherwise annoying or disrupting the victim's system or other systems (Mell *et al.* 2005, p. ES-1). Typical forms of malware include viruses, worms, Trojans, key loggers, rootkits and malicious mobile code.

Improving cybersecurity is not a straightforward problem. Notwithstanding rapidly growing investments in security measures, it has become clear that cybersecurity is a technological arms race that, for the immediate future, no one can win. Take spam, for instance. Several years ago, so-called open e-mail relays were a major source of spam. ISPs and other actors developed measures to collectively combat open relays, such as blacklisting. By the time adoption of these measures reached a critical mass, spammers had already shifted their tactics. As a result, the significant reduction in the number of open relays had hardly any impact on the amount of spam. More recently, the industry debated the use of Sender Policy Framework (SPF) as a way to combat the forging of the sender's mail addresses – a typical property of spam messages. While the industry was still discussing the merits of SPF, spammers were already successfully abusing SPF as a means to get even more messages past spam filters. The list of examples goes on and on.

While many would agree that cybersecurity needs to be strengthened, the effectiveness of many security measures is uncertain and contested. Furthermore, security measures may also impede innovation and productivity. Those involved in improving cybersecurity sometimes tend to overlook that the reason why the Internet is so susceptible to security threats – namely its openness – is also the reason why it has proven an enabling technology for an extraordinary wave of innovation and productivity growth. The benefits of the latter often outweigh the costs of the former – as in the case of online credit card transactions. From the start of moving their business online, credit card companies have struggled with rising fraud. This has not stopped them from expanding their online activities. The benefits of that growth were consistently higher than the associated costs of the increase in fraud. While growing in absolute terms, the level of online fraud in the United States has been dropping relative to the overall dollar amount of online transactions (Berner and Carter 2005). Rather than implementing far-reaching security measures that would restrict the ease of use of the system, credit card companies have adopted strategies to fight instances of fraud, up to the point where the costs of further reductions in fraud start to exceed the avoided damages.

All this means that total security is neither achievable nor desirable. In principle, actors need to make their own tradeoffs regarding what kind of security measures they deem appropriate and rational, given their business model. Clearly, business models vary widely for actors in the different niches of the complex

ecosystem surrounding information systems and networks – from ISPs at different tiers to software providers of varying applications to online merchants to public service organisations and to end users. All of these actors experience malware differently as well as the costs and benefits associated with alternative courses of action. In other words, many instances of what could be conceived as security failures are in fact the outcome of rational economic decisions, reflecting the costs and benefits perceived by the actors during the timeframe considered in those decisions.

What is needed then is a better understanding of these costs and benefits from the perspective of individual actors and of society at large. This report sets out to identify the incentives under which a variety of market players operate and to determine whether these incentives adequately reflect the costs and benefits of security for society – *i.e.* whether these markets generate externalities. It documents a research project designed with the goal of laying the groundwork for future policy decisions. We hope it supports OECD member countries in devising new policy options.

Research in the field of cybersecurity is undergoing a major paradigm shift. More and more researchers are adopting economic approaches to study cybersecurity, shifting emphasis away from technological causes and solutions. Most of this innovative research has yet to find its way into the realm of policy makers, let alone into the policies themselves. While reports like the OECD survey on the culture of security (OECD 2005) generally recognize that cybersecurity is more than a technological issue, the proposed measures are still mostly oriented in that direction: developing technological responses and efforts to stimulate their adoption. The technological responses are typically accompanied by legal efforts and intensified law enforcement.

Notwithstanding the necessity of these initiatives, they typically overlook the economic factors affecting cybersecurity – *i.e.* the underlying economic incentive structure. As Anderson and Moore (2006, p. 610) have argued, "over the past 6 years, people have realized that security failure is caused at least as often by bad incentives as by bad design." Many of the problems of information security can be explained more clearly and convincingly using the language of microeconomics: network effects, externalities, asymmetric information, moral hazard, adverse selection, liability dumping and the tragedy of the commons. Within this literature, designing incentives that stimulate efficient behaviour is central.

We can see the power of incentive structures around security threats everywhere. Take the distribution of viruses and other malware. During the second part of the 1990s, when the scale of virus distribution was rapidly increasing and countless end users (home, corporate, governmental) were affected, many ISPs argued that virus protection was the responsibility of the end users themselves. The computer was their property, after all. ISPs further argued that they could not scan the traffic coming through their e-mail servers, because that would invade the privacy of the end user. Mail messages were considered the property of the end users. About five years ago, this started to change, partly due to the growth of broadband and always-on connections. The distribution of viruses and worms had increased exponentially and now the infrastructure of the ISPs themselves was succumbing to the load, requiring potentially significant investment in network expansion. Facing these potential costs, ISPs radically shifted their position in response. Within a few years, the majority of them started to scan incoming e-mail traffic and deleting traffic identified as malignant as this had become a lower-cost solution than infrastructure expansion. *De facto* ISPs reinterpreted the various property rights associated with e-mail – *e.g.* regarding ownership of the message. Their changed policies have made e-mail based viruses dramatically less effective as an attack strategy.

Economics of information security and the OECD *Guidelines*

In 2002, the OECD released the *Guidelines for the Security of Information Systems and Networks* (OECD 2002a). A set of nine non-binding guidelines aim to promote "a culture of security" – that is, "a

focus on security in the development of information systems and networks, and the adoption of new ways of thinking and behaving when using and interacting within information systems and networks" – among "all participants in the new information society" (see Box 1). The guidelines reflect the shared understanding of OECD member countries as well as a variety of business and consumer organisations.

The "culture of security" that the guidelines aim to promote will be influenced by the incentive structures surrounding security tradeoffs. The focus on security may certainly be strengthened, but that in itself does not mean that actors will behave in ways that are beneficial to society. In other words, more attention to security does not equal better security decisions as long as economic incentives are ignored.

The next chapter provides a more detailed discussion of why this is the case. For now, it suffices to mention a few examples. Take the security investment levels of firms. Research has demonstrated that a focus on security may mean actively participating in information sharing with other firms. Under certain conditions, this actually leads to decreased investment levels. Also, a firm taking protective measures may create positive externalities for others – that is, benefits for others which are not reflected in the decision by that firm – which may reduce their investments to a level that is below the social optimum. Another example is the manufacturing of software. According to the *Guidelines* (OECD 2002b), "Suppliers of services and products should bring to market secure services and products." Even if it was clear what the term "secure software" means, many software markets do not reward such behaviour. Rather, they reward first movers – that is, those companies who are first in bringing a new product to market. This means it is more important to get to the market early, rather than first investing in better security. A final example relates to end users. The *Guidelines* argue that end users are responsible for their own system. In the case of malware, however, this responsibility may lead to security tradeoffs that are rational for the end users, but have negative effects on others. More and more malware actively seeks to reduce its impact on the infected host, so as not to be detected or removed, using the infected host to attack other systems instead of the host itself.

Box 1. OECD Guidelines for the Security of Information Systems and Networks

1)	Awareness Participants should be aware of the need for security of information systems and networks and what they can do to enhance security.
2)	Responsibility All participants are responsible for the security of information systems and networks.
3)	Response Participants should act in a timely and co-operative manner to prevent, detect and respond to security incidents.
4)	Ethics Participants should respect the legitimate interests of others.
5)	Democracy The security of information systems and networks should be compatible with essential values of a democratic society.
6)	Risk assessment Participants should conduct risk assessments.
7)	Security design and implementation Participants should incorporate security as an essential element of information systems and networks.
8)	Security management Participants should adopt a comprehensive approach to security management.
9)	Reassessment Participants should review and reassess the security of information systems and networks, and make appropriate modifications to security policies, practices, measures and procedures.

In short: the development of a "culture of security" is very sensitive to economic incentive structures. Whether such a culture will actually improve overall security performance requires a better understanding of the incentives under which actors operate as well as policies that address those situations where incentives produce outcomes that are not socially optimal. The project outlined in this report aims to contribute to this undertaking.

Report outline

An economic perspective on cybersecurity – and malware in particular – provides us with a more fruitful starting point for new governmental policies: incentive structures and market externalities. This report sets out to develop this perspective, building on the innovative research efforts of the past six years (for a brief overview of the existing literature, see Anderson and Moore 2007; Anderson *et al.* 2008). It is a first step in this direction but, given the complexity of the problem, more work will be needed.

Most of the research so far has been based on the methods of neo-classical and new institutional economics. While powerful, these methods are based on rather stringent assumptions about how actors behave – such as their rationality, their security tradeoffs and the kind of information they have – and how they interact with their institutional environment. Three key limitations of studies founded on these methodological assumptions are: *i)* they provide limited insight into how actors actually perceive the cost, benefits and incentives they face; *ii)* they have difficulties taking into account dynamic and learning effects, such as how a loss of reputation changes the incentives an actor experiences; and *iii)* they often treat issues of institutional design as rather trivial. That is to say, the literature assumes that its models indicate what market design is optimal, that this design can be brought into existence at will and that actors will behave according to the model's assumptions. If the past decade of economic reforms – including privatisation, liberalisation and deregulation – have taught us anything, it is that designing markets is highly complicated and sensitive to the specific context in which the market is to function. It cannot be based on formal theoretical models alone. Institutional design requires an in-depth empirical understanding of current institutional structures and their effects on outcomes. Even with such an understanding, it may not be possible to fully control the setup and working of a market as they are in part emerging from the interaction of multiple actors. However, it should be possible to nudge the system in the desired direction.

We propose to complement the existing research with qualitative field research. Only limited information as to how market players actually make their information security decisions is available in the public domain, which makes it difficult to calibrate any form of public policy. Our report presents our efforts to collect evidence on the security tradeoffs of market players, how they perceive the incentives under which they operate, which economic decisions these incentives support as well as the externalities that arise from these incentive structures. The objective of the report is to contribute to the debate on the economics of malware from an empirical and analytical perspective. It is not designed to explore and develop detailed policy recommendations.

Chapter II develops a framework to study the economics of malware. Both actors in the illegal and criminal world as well as actors within the information and communications sector respond to the economic incentives they face. After briefly exploring the connections between the markets for cybercrime and for cybersecurity, we focus on the latter. The economics of cybercrime is outside the scope of this study. The Chapter concludes by presenting a research design to qualitatively analyze the incentives of market players, their security decisions and the externalities that may arise within the market.

Chapter III reports the findings of the field work. Based on 41 interviews with 57 representatives of market players as well as governmental agencies and security experts, we discuss a variety of incentives

for Internet Service Providers, e-commerce companies (with a focus on financial service providers), software vendors, registrars and end users.

Chapter IV aggregates these findings and discusses the externalities that emerge from the incentives under which market players make security decisions. In some cases, externalities are borne by market players who are in a position to influence the security tradeoffs of the players from which the externality originate, bringing the value net as a whole closer to the optimum. In other cases, the externalities are borne by market players who cannot manage the originating security tradeoffs or they are borne by society at large. The report concludes with a summary discussion of the efficiency and distributional effects of externalities and an overall assessment of the costs of malware.

II. AN ECONOMIC PERSPECTIVE ON MALWARE

Information and communication technology (ICT) industries form a complex ecosystem and their services permeate most other economic activities. Security problems and the related economic costs to society may have two roots: *i)* they are the outcome of relentless attacks on the information and communication infrastructure by individuals and organisations pursuing illegal and criminal goals, and *ii)* given an overall external threat level, they may be aggravated by discrepancies between private and social costs and benefits which are the outcome of decentralised decision making in a highly interrelated ecosystem. Both actors in the illegal and criminal realms and within the information and communications system respond to the economic incentives they face.

In this complex value net (see Figure 1), economic decisions with regard to information security depend on the particular incentives perceived by each player. These incentives are rooted in economic, formal legal, and informal mechanisms, including the specific economic conditions of the market, the interdependence with other players, laws as well as tacit social norms. Within their own purview and constraints – for example, the available information may be incomplete – each player responds rationally to these incentives. It is critical for the economic efficiency of the whole value system that the incentives of the individual players are aligned with the overall conditions for social efficiency. In other words, the relevant incentives should assure that private costs and benefits of security decisions match the social costs and benefits. In cases of deviations between the private and socially optimal outcomes, the prevailing incentive mechanisms would ideally induce adjustments toward higher social efficiency.

Figure 1. Information industry value net

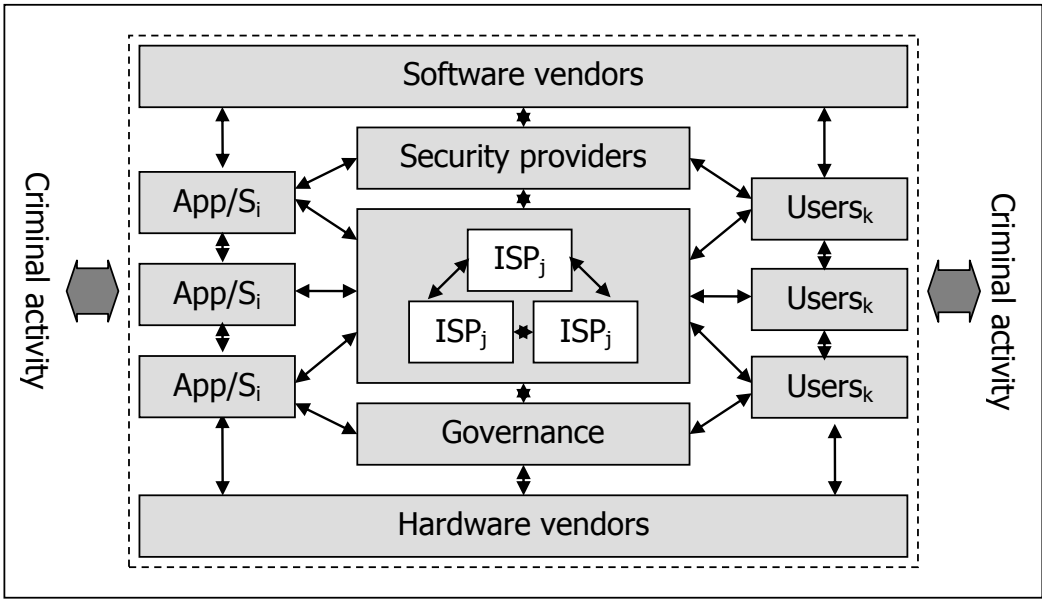

App/S_i … different types of application and service providers
ISP_j … different ISPs
$Users_k$ … different types of users (small, large, residential, business)

Misalignment between private and social efficiency conditions may take several forms. In case of incomplete information, the perceived incentives of individual players may deviate from the optimal incentives. A related issue is the problem of externalities, systematic deviations between the private benefits or costs and the social benefits or costs of decisions. Due to the high degree of interdependence, such deviations from optimal security decisions may cascade through the whole system as positive or negative externalities.

As the research on the economics of crime has illustrated, criminal activities may be analysed in a market framework. The activities in the market for cybercrime and cybersecurity are closely interrelated. Before the problem of incentives and externalities can be explored in more detail, we will, therefore, briefly explore the working of these markets and their linkages.

Cybercrime and information security

Figures 2 and 3 illustrate the interrelated nature of the markets for cybercrime and security. There are different ways to model the market for cybercrime. Becker (1968) and subsequent literature (see Ehrlich 1996; Becsi 1999 for overviews) suggest using a supply and demand framework to study criminal activity. Franklin *et al.* (2007) also employ an economic framework to study an underground economy based on "hacking for profit." We chose a slightly different representation than these studies, based on marginal analysis. It is reasonable to assume that a higher level of security violations is only possible at increasing cost. Furthermore, it is likely that the additional cost will increase more than proportionally as the extent of security violations increases.

On the other hand, the marginal benefits of additional security violations are a decreasing function of the level of violations. This is an expression of the fact that the most lucrative crimes will be committed first and that additional criminal activity will only yield lower marginal benefits. Criminals will extend their activities until the marginal cost of additional security violations approximates their marginal benefits. The magnitude of the benefits and costs of crime is dependent on a number of variables, some of which are affected by private and public measures to enhance security. A closer examination of these factors allows comparative assessments of market outcomes. It also sharpens understanding of the principal opportunities to intervene in the market to reduce cybercrime.

Technological change, the increased specialisation and sophistication in the production of malware, and the globalisation of the information and communication industries have all reduced the marginal cost of crime.[5] In turn, this cost decrease has dramatically expanded the supply of crime, as people from countries and regions with low opportunity cost of labour (which increase the net benefits of crime) join criminal activities. Such reduced marginal costs of security violations will shift the marginal cost of crime schedule downwards. Assuming that other things, especially the benefit relationship, remain unchanged, reductions in the marginal cost of crime will result in a higher level of security violations and vice versa.

Technological change and globalisation have also increased the benefits of crime. For example, the wider reliance on e-commerce and credit card transactions has increased the opportunities to exploit technical and personal security loopholes. The globalisation of the Internet has also enabled criminals to reach a larger number of potential victims. These changes shift the marginal benefit curve upwards (not captured in Figure 2). Other things being equal, this increase in the marginal benefits results in a higher

5 Statements as to the effect of changes in individual parameters or factors are typically made under the *ceteris paribus* assumption: that all other things remain equal. This is a widely used simplifying methodological tool to isolate changes in one or more variables in a highly complex interconnected system. Often, many factors will change simultaneously. A full grip on such changes will typically require some form of computer-based modelling or simulation.

level of security violations. The presence of both effects explains much of the increased level of activity of security violations. In principle, however, opposite shifts of the marginal cost and benefit curves may be achieved by appropriate measures.

Figure 2. Markets for crime and security

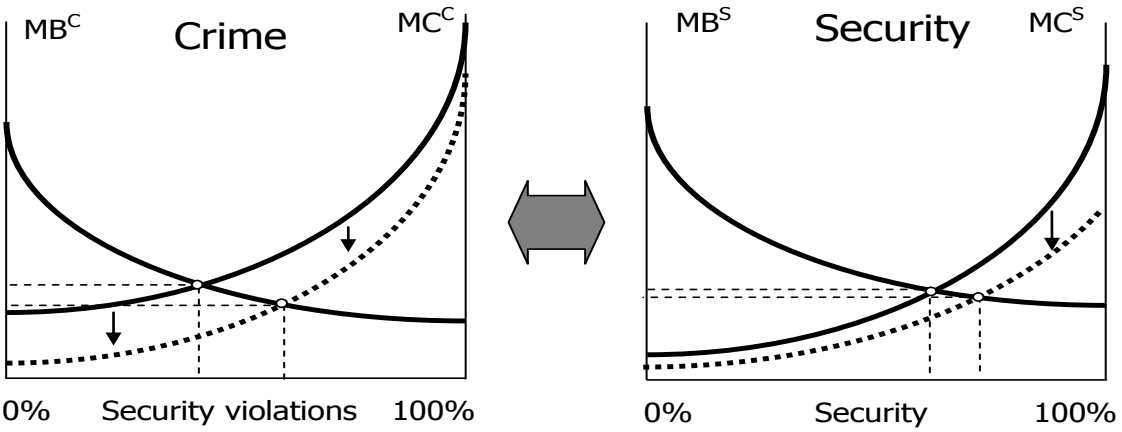

Figure 3.

MB^C … marginal benefits of crime
MC^C … marginal costs of crime
MB^S … marginal benefits of security
MC^S … marginal costs of security

The market for security can be analysed using a similar approach. It is reasonable to assume that higher levels of security can only be achieved at higher marginal costs. On the other hand, the marginal benefits of security will decrease. Unless the benefits exceed the cost throughout, the resulting optimal level of security will be below 100%, at least on an aggregate level.[6] Changes in the costs of providing security and the benefits of having security will shift the marginal cost and benefit schedules and affect the market outcome. A reduction in the cost of security, for example, due to the availability of more efficient and cheaper filtering software or a new network architecture that might reduce the propagation of malware, will, *ceteris paribus*, result in a higher level of security. Likewise, higher benefits of security, perhaps because of the utilisation of more mission-critical applications, will, other things being equal, result in a higher level of security. However, such initial changes may result in subsequent adjustments by other actors, who might reduce their expenditure for security in response, leaving the overall effects on the resulting security level ambiguous at best (see the arguments in Kunreuther and Heal 2003).

6 It is possible that for some services and applications 100% security levels are required (hence the benefits higher than the cost, even at a level of 100% security) and that the requisite cost will be incurred. It is unlikely, though, that this will hold for all services and applications.

Figure 4. Markets for crime and security

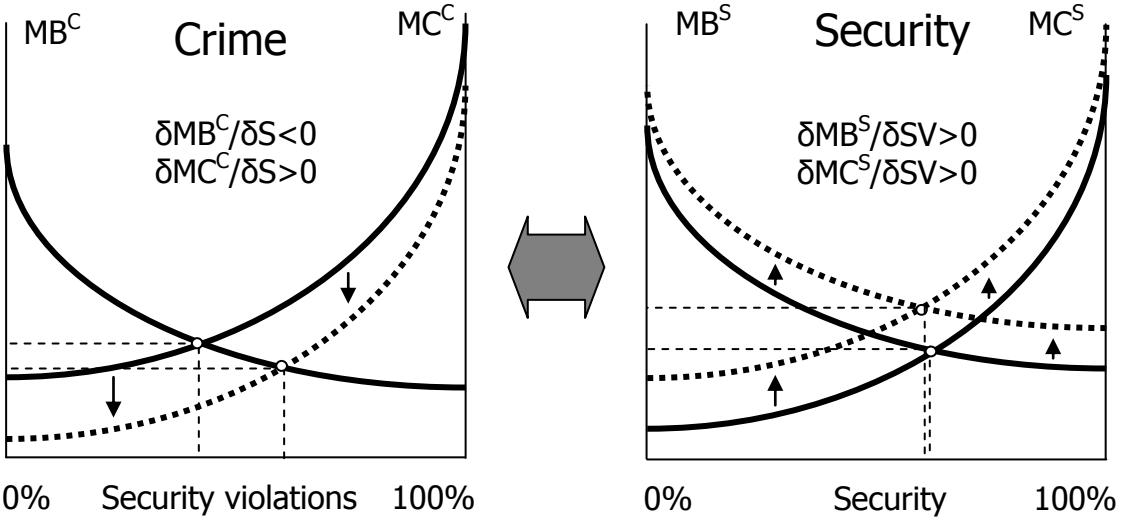

MB^C … marginal benefits of crime
MC^C … marginal costs of crime
MB^S … marginal benefits of security
MC^S … marginal costs of security
$\delta MBC/\delta S<0$ expresses the changes of the MBC curve in response to a change in the level of security S. The negative sign implies that the marginal benefits of crime move in the opposite direction from marginal changes in security, *i.e.* increased security reduces the marginal benefits of crime, all other things being equal.

The markets for cybercrime and security are highly interrelated (Figure 3). Activities in the market for cybercrime affect the market for security and vice versa. Most likely, an increased level of security violations will increase the marginal benefits and the marginal costs of security, shifting both schedules upwards. On the contrary, a lower level of security violations resulting from the market for crime will shift both schedules down. On the other hand, variations in security will have corresponding effects on the market for crime. Increased security will increase the marginal cost of security violations and it will reduce the marginal benefits of crime.[7] The net impact on the overall level of security is difficult to predict and will depend on the relative strengths of variations in security violations on the costs and benefits of security. A higher level of security violations could result in a lower level of security, an unchanged level of security, or even a higher level of security. Without any specific policy intervention, the interaction between the two markets may resemble an arms race.

There is an asymmetry in the effects of each market on the other. On the one hand, an increased level of security violations may or may not affect the level of security. However, for all actors it will likely result in higher costs of maintaining a certain level of security. On the other hand, a higher level of security will induce changes in the market for crime in that it will increase the marginal cost of security violations and, at the same time, reduce the marginal benefits of crime. Both effects will mutually reinforce each other, thus contributing to a lower level of security violations. As parameters in each of the markets change continuously, the outcomes of the resulting dynamic mutual adjustment are difficult if not impossible to model, although the directions of change seem to be robust.

[7] More formally, the partial derivatives can be expressed as: $\delta MBC/\delta S<0$, $\delta MCC/\delta S>0$, $\delta MBS/\delta SV>0$, $\delta MCS/\delta SV>0$.

This framework also gives first, high-level insights into the measures that are available to influence the overall outcomes. Such measures can target the market for cybercrime and/or the market for security. Measures such as increasing the cost of cybercrime by increasing the associated penalties, strengthening national and international law enforcement, and increasing the difficulty of registering and maintaining fraudulent domains and websites will affect the market for crime directly and also have repercussions on the market for security. Most likely such measures will reduce the overall level of security-related costs. For reasons discussed above, it is less certain that such measures will increase the level of security, as accepting a certain level of insecurity is economically rational.

Measures affecting the overall incentive compatibility in the security markets range from forms of industry self-regulation to forms of co-regulation and government intervention. They encompass a wide spectrum of measures such as requiring that security features are enabled by default, recommendations to ISPs to adopt best practices with regard to security on their networks, information campaigns to alert users to security risks, and changes in the ways domain names are registered. None of these measures is a panacea but they help better align individual incentives with social efficiency requirements.

Incentives and economic decisions

Economic incentives are the factors that influence decisions by individuals and individuals in organisations. A close examination of the incentives of the stakeholders in the information industry value network to undertake measures to prevent or mitigate the costs associated with malware is thus critical to a full understanding of the economics of malware. Such actions include investment in security, investment in technical means to prevent or at least control problems caused by malware, response sequences in case an intrusion has happened or an attack is unfolding. The relevant sets of incentives are most likely different for each stakeholder. Hence we attempted to get a detailed account of the perceived incentives from experts in the respective segments along the value net. Moreover, the incentives may complement each other, they may form a trade-off, or they may even work at cross-purposes. An important goal of our analysis was, therefore, to examine the aggregate interaction of the individual incentives faced by stakeholders at the sector level. As systems of incentives have many feedback loops, it is typically very difficult to determine the net effect of a system of incentives. At this stage of the project we used a qualitative assessment approach.

Economic incentives shape decisions in for-profit commercial firms, non-profit social groups, public and private sector governance institutions, as well as not-for-profit forms of production and collaboration. Incentives are often classified in monetary (remunerative, financial) and non-monetary (non-financial, moral) factors. Financial incentives include factors such as tying the salary of an employee to corporate performance, the ability to make a super-normal profit by pursuing a risky innovation, or the bottom line effects of potential damage to a firm's reputation. Non-financial incentives encompass norms and values, typically shared with peers, and result in a common understanding as to the right course of action or the set of possible actions that should be avoided in a particular situation. Financial incentives typically connect degrees of achievement of an objective with monetary payments. Non-financial incentives work through self-esteem (or guilt) and community recognition (or condemnation).

In practical decision making, incentives can be seen as the motives for selecting a specific action or the rationales for preferring one course of action over another. As the discussion of reputation effects illustrates, it is sometimes necessary to distinguish between short-term and long-term effects. Characteristic features describing incentives are their power (low-powered to high-powered) and directionality (positive or negative relation to goals of decision).[8] An important question is the relation

[8] Mechanisms operating towards improving an objective are typically referred to as "incentives" whereas those operating in the opposite direction are referred to as "disincentives."

between the structure and power of the relevant incentives and the objectives of decisions. The full set of incentives at work typically consists of a bundle of specific, more narrowly defined, incentive mechanisms. These incentive mechanisms may work in the same direction or conflict with each other. If feedback loops between incentives exist it is often difficult to determine their overall net effect. However, it is possible to establish the effect of a single incentive mechanism under the methodological assumption that all other factors remain constant (*ceteris paribus*). For example, for software vendors the reputation mechanism *ceteris paribus* works toward increased information security but potential first mover advantages in information industries may, *ceteris paribus*, lower the incentives to invest in information security.

Incentive-compatibility refers to a situation in which an incentive is structured in a way so as to contribute to the stated goals of an individual or an organisation. To assess incentive compatibility, the direct and indirect links between an incentive mechanism and the objective being pursued will have to be examined. Incentive compatibility may exist at the level of a single incentive mechanism, the bundle of incentives at work for a specific stakeholder, or the entire sector under consideration. Given the potential for trade-offs and even direct conflicts between incentives, incentive compatibility is much more difficult to ascertain at the level of stakeholders and the industry at large. It is a particular challenge in an industry as highly inter-related as advanced information and communication industries are. To be affected by an incentive mechanism, individuals need to be cognizant of its existence, its directionality, and its power. Incentives that exist on paper but are ignored by the decision makers must either be seen as zero-powered or as irrelevant incentives. Therefore, it is possible to reveal the existing incentive structures of the stakeholders in the information value net by asking experts and decision makers for an in-depth account.

Externalities

Externalities are forms of interdependence between agents that are not reflected in market transactions (payments, compensation). Which phenomena are identified as externalities depends to a certain degree on the specification of legal rights and obligations in the status quo. If these rights and obligations are only vaguely defined they may need clarification by legislatures, courts and in private contractual agreements.[9] If such clarification is afflicted with transaction costs, rational individual actors affected by the externalities will not internalise them if these costs exceed the potential benefits of internalisation. In this case, only a collective actor (*e.g.* a business association, government) may be able to address these uncompensated externalities.

In the formulation of the mainstream economic model, these interdependencies lead to deviations from a socially optimal allocation of resources. Negative externalities result in an overuse or overproduction compared to the social optimum whereas positive externalities lead to an underuse or underproduction of the resource afflicted with the externality (Friedman 2002, pp. 599). External effects are often classified according to the agents that are involved. Frequently, producers and consumers are distinguished, yielding a two-by-two matrix of producer to producer, producer to consumer, consumer to producer and consumer to consumer externalities (Just *et al.* 2004, pp. 527).

An alternative typology distinguishes between technological and monetary externalities (Nowotny 1987, p. 33). Technological externalities are said to exist if, at constant product and factor prices, the activities of one agent directly affect the activities of another. Pecuniary externalities exist, if the activities of one agent affect the prices that need to be paid (or may be realized) of other agents. Early contributions to the subject, for example, by Marshall (1920) or Pigou (1932), treated externalities as an exception, a rare

[9] This seems currently the case in many countries. See for example: 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 http://www.bsi.de/literat/studien/recht/Gutachten.pdf.

anomaly in a market system. However, the increasing concern with environmental issues since the 1960s made clear that such interdependencies are pervasive and part and parcel of real world market systems.

This is particularly true for information and communication networks, which raise several new and unique issues. The high degree of interconnectedness amplifies the interdependencies between participants in the network. Both negative and positive effects that are not reflected in market transactions may percolate widely and swiftly through electronic communication networks. In some types of networks, such as peer-to-peer arrangements, agents take on dual roles as consumers as well as producers of information and other services. Many users of cyberspace view it as a commons, in which transactions take place according to a gift rather than marketplace logic. Moreover, often, for example, in the case of Trojans, externalities are generated without the explicit consent or knowledge of an individual user. All these factors influence the prevalence of externalities and complicate possible ways to address them.

Origins of externalities in networked computer environments

External effects may originate at different stages of the value net in networked computer environments. Depending on the origin of the externality, the individual decision-making calculus causing the externality may be different. In any case decision makers focus on costs and benefits relevant to the individual agent and neglect costs or benefits of third parties.[10]

Table 1 provides an overview of the sources and forms of externalities in networked computer environment. The table captures the main stakeholders, but not necessarily all of them. Agents in the column are the sources of externalities whereas agents in the rows are the recipients. Not all agents may cause externalities on all others and some of the effects may be more likely or stronger than others. By definition, an agent cannot exert an externality on itself, although it may create an externality for another agent in the same category. For example, the lax security policy of one ISP may create externalities for other ISPs.

A first source of possible externalities is software vendors. When deciding the level of investment in activities that reduce vulnerabilities, software vendors will only take their private costs and benefits into account (Schneier 2000). Sales of software are dependent on the reputation of the firm. If this reputation effect is strong, the firm will also be concerned about the security situation of the software users. However, it is likely that such reputation effects are insufficient to fully internalise externalities. This situation is aggravated by the unique economics of information markets with their high fixed costs and low incremental costs, the existence of network effects which create first-mover advantages, and the prevalence of various forms of switching costs and lock-in. These characteristics provide an incentive for suppliers to rush new software to the market (Anderson 2001; 2002; Shostack 2005). They may also lead to the dominance of one or a few firms, increasing overall vulnerability due to a "monoculture" effect (Böhme 2005).

[10] In a dynamic context, reputation effects may mitigate some of the externalities, see the discussion below.

Table 1. Origins and forms of externalities in networked computer environments, as seen from the source of the externality

	Software vendors	ISPs	Large firms	SMEs	Individual users	Criminals
Software vendors	Level of trust, reputation	Risk of malevolent traffic	Level of software vulnerability	Level of software vulnerability	Level of software vulnerability	Hacking opportunities
ISPs	Level of trust, reputation	Volume of malevolent traffic	Risk of proliferating attack	Risk of proliferating attack	Risk of proliferating attack	Hacking opportunities
Large firms	Level of trust, reputation	Volume of malevolent traffic	Risk of hosting or proliferating attack	Risk of hosting or proliferating attack	Risk of hosting or proliferating attack	Hacking opportunities
SMEs	Level of trust, reputation	Volume of malevolent traffic	Risk of hosting or proliferating attack	Risk of hosting or proliferating attack	Risk of hosting or proliferating attack	Hacking opportunities
Individual users	Level of trust, reputation	Volume of malevolent traffic	Risk of hosting attack	Risk of hosting attack	Risk of hosting attack	Hacking opportunities
Criminals	Level of trust, reputation	Resource use, reputation	Resource use, Costs of crime	Resource use, Costs of crime	Resource use, Costs of crime	Hacking opportunities

Source: own construction.

Whether they be large corporate users or small and medium-sized firms, security investments by firms to reduce vulnerabilities are likewise afflicted with externalities, as discussed by several authors (Gordon and Loeb 2002; Vijayan 2003; Camp and Wolfram 2004; Schechter 2004; Chen *et al*. 2005; Rowe and Gallaher 2006). Profit-maximizing firms, all other things being equal, will attempt to invest in information security until the (discounted) incremental private benefits of enhanced security are equal to the (discounted) costs of that investment. A firm will therefore not invest until the security risk is fully eliminated but only as long as the expected costs of the threat are higher than the cost of increasing information security. Costs that the firm imposes on third parties will not be considered in this calculus (unless they indirectly affect a firm's decision making, for example, because of reputation effects).

Likewise, benefits that a security investment bestows on third parties will also not be reflected in this decision. Under conditions of imperfect information and bounded rationality, firms may not be able to determine this private optimum with precision but they will try to approximate it. In any case, neither the negative external effects of investments falling short of the social optimum nor the positive externalities of investments that go beyond that optimum are taken into consideration. Individual firm decisions may thus systematically deviate from a social optimum that takes these interdependencies into account.

Individual users are seen by many as one of the weakest links in the value chain of networked computing (Camp 2006). Larger business users often consider their decisions in an explicit cost-benefit framework. In contrast, small business and individual users often do not apply such instrumental rationality (LaRose *et al.* 2005; Rifon *et al.* 2005). Nevertheless, when making decisions as to security levels, they consider their own costs and benefits (but not those of other users). Individual users are particularly susceptible to non-intrusive forms of malware, which do not use up significant resources on the user end (*e.g.* computing power, bandwidth) but create significant damage to other machines. Consequently, the risk of attack for all other users and the traffic volume on networks is increased causing direct and indirect costs for third parties.

ISPs may inflict externalities on other agents in the value chain as well as on each other. Some malware may increase traffic and hence ISP costs only incrementally. In this case, the ISP may have little

incentive to incur additional costs to engage in traffic monitoring and filtering. Even if users cause significant traffic increases, an ISP with a lot of spare capacity may not see anything but very incremental cost increases, again limiting the incentive to invest in security upgrades to reduce malware-related traffic.

Information security externalities appear in several forms, including direct costs or benefits and indirect costs and benefits. Direct costs include damage caused to other stakeholders (such as corrupted data or websites, system downtimes) and the cost of increased preventative security expenses by other stakeholders (including cost of software and security personnel). Indirect costs include reduced trust within computer networks (for example, if nodes maintain lists of trusted other systems) and of users in information networks, the ability of hackers to increase the effectiveness of attacks by subverting more machines, and the ability of hackers to hide their traces (Camp and Wolfram 2004). They also include the potentially high costs associated with the reduced willingness of consumers to engage in e-commerce.

Externalities in a dynamic framework

In networked computer environments with rapid technological change, externalities need to be understood in a dynamic framework. Most importantly, learning and reputation effects need to be considered. Reputation and learning may happen at different time scales and with different intensity in the various components of the value net. They will also differ within markets, for example enterprise market software as opposed to mass market software. In any case, they may counteract and reduce the magnitude of negative externalities and possibly enhance positive externalities. Moreover, the activities of firms to disclose vulnerabilities will influence the magnitude of externalities.

Table 2. Externalities with reputation

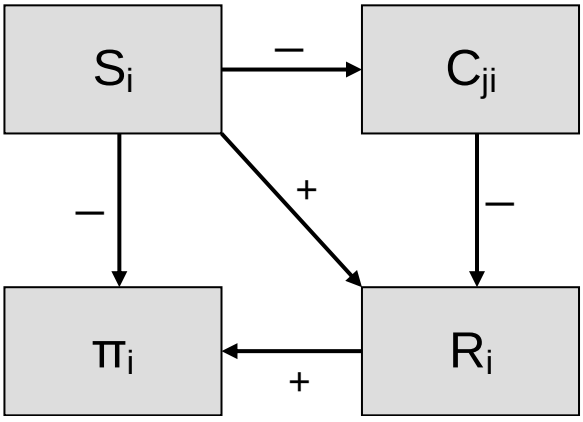

S_i security investment of firm i
C_{ji} cost for firm j cause by sub-optimal security investment by firm i
R_i reputation of firm i
π_i profits of firm i

Table 3 illustrates the reputation effect for the case of a software vendor (plus and minus signs indicate whether the two variables move in the same or the opposite direction). Other things being equal lower expenses for system testing and refinement by firm i (S_i) will reduce sunk costs and hence increase the profits (π_i) of the firm. However, costs may be externalised onto other firms, indexed j (C_{ji}). If these costs affect the reputation of firm i (R_i), profits may be reduced, especially if the reputation effect works swiftly. In this case, at least part of the potential externality is internalised and the deviation between private and social optimum is reduced. One form of strengthening the reputation mechanism is trusted-

party certification. As Edelman (2006) and Anderson (2001) point out, given present liability rules, these firms face an adverse selection incentive in that they do not face any consequences for issuing wrong certificates.

In a dynamic perspective, the incentives to disclose vulnerabilities need to be considered (Cavusoglu *et al.* 2005). Disclosure exerts a positive externality (Gal-Or and Ghose 2003; Gal-Or and Ghose 2005) onto other stakeholders. Under certain conditions, disclosure incentives may be sufficiently strong to shrink the conditions under which deviations between the private and social optimum occur to a minimum (Choi *et al.* 2005) .

Research design

Our evaluation started with an exploration of the incentives at work in the individual organisation and those related to the decisions of other competing or complementary organisations. The reliability of the information is increased if interdependent stakeholders present compatible pictures of the relevant incentives and their effects. Attempts were made to interview several organisations in each segment of the value chain to develop narratives that are as coherent as possible. In a subsequent analytical step, these individual narratives were then integrated to assess the overall incentive structure of the sector and the resulting externalities.

Data collection

In the course of 2007, we conducted 41 in-depth interviews with 57 professionals from organisations participating in networked computer environments that are confronted with malware. Firms from the following components of the value net were approached:

- Internet Service Providers
- E-commerce companies, including online financial services
- Software vendors
- Hardware vendors
- Registrars
- Security service providers
- Different types of end users
- Governance institutions (regulators, consumer protection agencies, CERTs)

A full list of respondents can be found in the Appendix. Our empirical effort extends the preliminary work on firms and end users (*e.g.* Dynes *et al.* 2005; Camp 2006; Dynes *et al.* 2006; Poindexter *et al.* 2006; *e.g.* Rowe and Gallaher 2006).

The interviews were carried out using a semi-structured questionnaire, adapted for the specific situation of the interviewee. In each instance, we asked how the organisation was confronted with malware, what its responses were, what tradeoffs were associated with these responses, and how the organisation was affected by the actions of other market players. As is common practice in the social sciences, we have treated all interview data as confidential, so as to enable the interviewees to share information with us as freely as possible. Consequently, no interviewee or organisation is identified by name in relation to specific data and all quotes have been approved by the respective individuals/organisations beforehand for publication. All statements in the report are based on interview transcripts and other documents supporting the findings. Although this limits the direct verifiability from

readily available public sources, we felt that given the exploratory stage of research in this area, our approach would enable us to get better insights into market-sensitive economic data and decision making.

Scope and limitations

Before turning to the empirical findings, it is important to note the scope and limitations of this study. The global and heterogeneous nature of the ecosystem of Internet services implies that any study of incentives is almost by necessity an exploratory study. The limited time and budget available for this study allowed for a limited number of interviews in six countries. The majority of the interviews took place in the United States and the Netherlands, with additional interviews in the United Kingdom, France, Germany and Australia. The next section presents our findings for five of the market players we interviewed:

- Internet Service Providers
- E-commerce companies, including online financial services
- Software vendors
- Registrars
- End users

We intended to also describe the incentives for hardware vendors, but we were unable to secure sufficient interviews with hardware vendors to provide the basis for such a description. The examination of the incentives was based not only on interviews with the market players themselves, but also on conversations with people who have expertise on the current threats and governance of the ecosystem of information services, such as regulators, CSIRTs, ICANN, security services providers, and researchers.

While these interviews have proven to be highly informative, the findings drawn from them should be read with caution. First of all, it is reasonable to assume that the set of interviewees is influenced by some degree of self-selection. ISPs, for example, are more likely to respond favourably to an interview request about the economics of malware if they have security policies in place that are at least on par with other ISPs, if not better. That said, some of the organisations we interviewed are publicly known for a less than stellar track record with regard to security – which they often explicitly acknowledged during the conversations. Second, the empirical findings report on how stakeholders themselves describe what they are doing and why. In other words, we report on the perceptions of the interviewees, not some independent assessment of their actions and the factors driving them. Whenever possible, we did cross-check information provided to us against the information from other interviews and against publicly available data, such as security reports, surveys and research publications. Third, the interviews touch on many issues that concern proprietary or otherwise confidential data. Interviewees were not always able to share this data with us and if they were, we were constrained in reporting them. Fourth, and last, our interviews involved six different legal jurisdictions. Some incentive mechanisms are generic but others are context-specific. Our approach hence provided us with a sense of the degree to which certain findings were country-specific and therefore could not fully reflect the heterogeneity of all OECD members.

These circumstances make it more difficult to generalise our findings. However, very little empirical field work has been done in this area so far. In light of the rapidly increasing political attention given to the issue of malware and the policy initiatives currently under debate, this is a critical omission. Our study contributes to overcoming this omission. At the very least, it makes clear the urgency of developing a further-improved in-depth understanding of the economics of malware to increase the probability of policy interventions to succeed.

III. SECURITY DECISIONS AND INCENTIVES FOR MARKET PLAYERS

Players across the ecosystem of information services are confronted with malware in different ways. Their responses to malware are motivated by the specific incentives under which they operate. To better understand these incentives and their effects, a qualitative field research project was designed. In the course of 2007, the research team conducted 41 interviews with 57 respondents from a broad cross section of organisations – for more information on the research design, see Chapter II; for a full list of interviewees, see Appendix. Below, we discuss our findings on the security-related incentives of Internet Service Providers (ISPs), e-commerce companies (with a focus on online financial services), software vendors, registrars and end users in more detail. Interviews were also conducted with representatives of organisations governing security issues (such as CERTs, regulatory agencies), representatives from security service providers, and other researchers.

Internet service providers

Over the past years, it has proven extremely difficult to improve the security of end users. Reliable estimates are hard to come by, but several of our sources subscribed to estimates available elsewhere that 20-25% of computers worldwide are at risk because their owners are unwilling or unable to adequately secure their systems (*e.g.* BBC News 2007; *e.g.* House of Lords 2007a, p. 29; Weber 2007). Other estimates are considerably lower – *e.g.* Trend Micro published a figure of 7% (Higgins 2007b). Nevertheless, even these estimates still imply tens of millions of compromised machines. Given the enduring problems around end user security and its effects on the wider network, it seems inevitable that attention would shift to other players in the ecosystem. The role of ISPs in improving Internet security has been a particular focus of recent debates.

While the term ISP is used to cover a variety of businesses, typically ISPs are defined as providers that offer individuals and organisations access to the Internet. Many ISPs offer related services to their customers, which is why the term sometimes refers to hosting providers and content providers. We have focused our analysis primarily on ISPs as access providers.

What incentives do ISPs have to reduce the problem of malware? One view is: very few, if any. Recently, the UK House of Lords Science and Technology Committee published a report which states: "At the moment, although ISPs could easily disconnect infected machines from their networks, there is no incentive for them to do so. Indeed, there is a disincentive, since customers, once disconnected, are likely to call help-lines and take up the time of call-centre staff, imposing additional costs on the ISP." (House of Lords 2007a, p. 30)

ISPs may unwittingly reinforce the impression that they have few if any incentives to improve the security of their services. During the inquiry that led to the House of Lords report, ISPs argued that the current approach to self-regulation should not be changed. The resistance of most ISPs to increased government involvement led the committee to conclude that the ISPs were simply maintaining the status quo, rather than reducing the problem. The latter, however, does not follow from the former. The resistance to government involvement does not mean that ISPs are not increasing their efforts to fight malware. In fact, the committee itself also cites evidence from an ISP who in fact disconnects customers whose machines had been infected and then helps them back online. A survey from the EU's European Network and Information Security Agency found that 75% of ISPs report that they quarantine infected machines

(ENISA 2006). This figure does not include any indication of the scale at which ISPs are quarantining infected machines – a point to which we return in a moment. The evidence does, however, clearly question the earlier statement by the committee – and others – that ISPs have no incentives to disconnect infected machines. Either the statement is wrong, or ISPs are assumed to behave irrationally. Our evidence suggests the former.

All ISPs we interviewed described substantial efforts in the fight against malware, even though they are operating in highly competitive markets and there is no governmental regulation requiring them to do so. All of them were taking measures that were unheard of only a few years ago. Most of the interviewees dated this change to around 2003, when it became obvious that it was in the ISPs own interest to deal with end user insecurity, even though formally it was not their responsibility. Several incentives help explain why the ISPs see these efforts as being in their own interest.

Costs of customer support and abuse management

An understanding of these incentives could start with this statement by a security officer of a smaller ISP: "The main [security-related] cost for ISPs is customer calls." The same view was expressed with minor variations by several other interviewees. A medium-sized ISP told us that an incoming call to their customer center costs them EUR 8 on average, while an outgoing call – for example, to contact the customer regarding an infected machine – costs them EUR 16. The costs for e-mail were similar. When we mentioned these numbers during subsequent interviews with other ISPs, they confirmed that their costs were in the same range. The incentive here is that security incidents generate customer calls, thus quickly driving up the costs of customer care. The ISPs may not be formally responsible for the customers' machines; in reality many customers call their ISP whenever there is a problem with their Internet access. Regardless of the subsequent response of the ISP, these calls increase their costs. An interviewee at a large ISP told us that their customer support desk was a substantial cost for the company and that the number of calls was driven up by infections of their customers' machines. He further added that almost all of their outgoing security-related calls had to do with malware.

Of course, many forms of malware do not manifest themselves explicitly to customers. Nevertheless, as security problems rarely come alone, lax security generally tends to increase customers calls. Furthermore, even if customers have not noticed anything wrong, their compromised machines may generate abuse notifications to their ISP from other ISPs who monitor incoming spam or malware from the customer's IP address. Similar to customer contact, dealing with abuse notifications drives up costs because it requires trained staff. Tolerating more abuse on the network raises the number of notifications that have to be investigated, responded to and acted upon. Acting may mean filtering the customer's connection or even suspending it altogether, until the problem gets resolved. All the ISPs we interviewed have procedures in place for handling abuse notifications and do in fact filter and suspend connections, though with varying frequency. All of them also mentioned a small number of cases where extreme forms of abuse led to the termination of the contract.

Abuse notifications can come through different channels, most notably through e-mail sent to the abuse desk – typically abuse@provider.com – and through the informal networks of trusted security professionals that exist across ISPs, CSIRTs and related organisations. The latter carry more weight, as they come from known and trusted sources, but all have to be dealt with in some form. Many of these notifications are automated. Several ISPs reported using the so-called AOL Feedback Loop, which sends notifications of any e-mails that are reported as spam by AOL recipients back to the administrator of the originating IP address. As with customer complaints, not all malware infections will result in abuse notifications. One ISP reported internal research into the degree to which notifications adequately represented the size of the security problems on their networks. They found that only a small percentage of the compromised machines they saw on their network showed up in the notifications. Still, ISPs notifying

each other of security problems is an important mechanism. In fact, in some cases, they are critical. In some European countries ISPs have interpreted the stringent privacy regulations in ways that substantially limit their ability to monitor their own network. In these cases. they rely heavily on notifications coming in from other ISPs, which then allow them to initiate their own investigation.

For the ISPs we interviewed, customer contact and abuse notifications are a strong incentive to invest in security both at the network level, as well as at the level of the customer. One medium-sized ISP estimated they were spending 1-2 % of their overall revenue on security-related customer support and abuse management. This also helps to understand why more and more ISPs are offering "free" security software or "free" filtering of e-mail – that is, the costs of these services are included in the subscription rate. One ISP described how about four years ago they started offering virus filters for e-mail as a paid service, but soon thereafter decided to provide them for 'free': "After six months, all ISPs [offered these paid security services], so it was no longer a unique selling point. Plus, we could not get more than 10 % of our customers to buy the service... We did not actually do the math, but we figured that by offering it to all our customers within the current rate, we would be better off.... We already paid the AV license. If people have the option to pay for it or not to pay for it, they do not."

There is another way of responding to these incentives, however: Don't respond to abuse notifications and avoid customer contact altogether. A class of ISPs is doing exactly this. What is stopping other ISPs, including the ones we interviewed, from doing the same? Here, we came across two interrelated relevant incentives: blacklisting and brand damage.

Costs of blacklisting

Blacklisting is a loosely used term typically referring to the practice of ISPs of using so-called DNS Blacklists (DNSBL) to filter incoming traffic. Mail servers, for example, may be configured to refuse mail coming from IP addresses, IP ranges or whole networks listed on a DNSBL. There is a wide variety of blacklists available and ISPs may use them in different combinations. According to many interviewees, most ISPs use blacklists nowadays. Most of the lists are free and run by volunteers, though their operations may be funded through external sources. Each DNSBL has its own criteria for including an IP address in the list and its own procedure for getting an address off the list. Spamhaus, an international non-profit organisation funded through sponsors and donations, maintains several famous blacklists – though they prefer the term block lists – which they claim are used to protect over 600 million inboxes. One of their lists contains the addresses of "spam-sources, including spammers, spam gangs, spam operations and spam support services"; another list focuses on botnets which run open proxies. It should be noted at this point that blacklisting, while potentially powerful, has drawn its own criticisms – regarding, among other things, vigilantism of blacklist operators, listing false positives, the collateral damage that may come with blacklisting certain IP addresses or ranges, and the financial motives of some list operators. Furthermore, blacklists have suffered from legal threats, where spammers on occasion were successful in obtaining court verdicts against being blacklisted (*e.g.* Bangeman 2006; Heidrich 2007). Within this report we focus on how blacklisting works as an incentive for ISPs.

Blacklisting provides an incentive to invest in security because it ties in with the incentives mentioned earlier. One interviewee at a medium-sized ISP told us about a security incident where 419 spammers set up over a thousand e-mail accounts within their domain and then started pumping out spam. That got the ISP's outbound mail servers blacklisted, which resulted in 30 000 calls to their customer center by customers who noticed their e-mail was no longer being delivered. That number does not include the incoming abuse notifications, of which there were "even more." After this incident, the company changed the procedure through which new customers can set up e-mail accounts; they invested millions in equipment to monitor their network; and they started blocking port 25. "It took us years to get a procedure approved to be able to block port 25. It costs nothing. But the business units did not want us to be able to

shut it down, because of their clients. They now understand that it is in the interest of their clients, to avoid blacklisting."

Blacklisting directly impacts the ISP's business model. A security officer at a large ISP explained that being blacklisted led to a much more proactive approach to remove bots from the network, including the purchase of equipment that automates the process of identifying infected machines on the network. The ISP contacts around 50 customers per day and, if the customer does not resolve the problem, the connection is suspended. When asked how they got the business side of the company to approve this policy, he answered: "They hated it at first. But at the end of the day, the media fallout by being cut off by AOL and MSN is too big. The big ISPs, they use very aggressive [DNSBL] listings. They take out whole IP ranges. We used to be hit hard and entire ranges of our IP addresses were blacklisted."

There are various levels of blacklisting used to incite a response from an ISP. At the lower end, we find blacklisting of individual IP addresses, *i.e.* an individual customer. This has "exactly zero impact on the ISP," said a security expert. Only when they start to accumulate, might they get the ISP's attention. The expert explained that ISPs mostly ignore listed individual IP addresses, because of the costs of dealing with them – *e.g.* customer support – and because the IP addresses gets taken off of the blacklist as spammers or attackers move on to other infected machines. After a few months, the level of active infected machines on the ISP's network might be equally high, but it is a different set of individual IP addresses that are now blacklisted. Blacklisting IP ranges and blacklisting outbound mail servers are a more powerful incentive. These typically do get the ISPs attention and lead to remedial action on their end, although it varies whether or not the ISP remains vigilant. The most extreme form is blacklisting an entire network, *i.e.* all IP addresses of an ISP. This is only used against semi-legitimate ISPs who do not act against spam and against known spam-havens.

Costs of brand damage and reputation effects

The "media fallout" mentioned by the interviewee points to a more general concern with brand damage that was mentioned by many interviewees as an incentive to invest in security. With few exceptions, these ISPs want to present themselves as responsible businesses (Arbor Networks 2007) providing safe services for their customers. A related incentive is the reputational benefits of offering security services. The increasing attention on Internet security – or rather, to the lack thereof – is creating demand for such services. One interviewee said: "The banks ask us for 'clean pipes.' We do not know what that means exactly, but they ask us anyway. We're looking into what we can do for them." The past years have witnessed the emergence of managed security service providers, either by conventional ISPs taking on security services, by security providers adding Internet access or by new businesses altogether.

It is unclear how strong this incentive is. For the large and medium-size business market, the ISP's image in terms of security may be a significant factor. For the consumer market, many interviewees argued that customers care about price first and foremost and, thus, Internet access is marketed primarily on price. Furthermore, even if they do care about security, most customers will find it very difficult to assess the security performance of one ISP relative to its competitors. Nevertheless, the more significant finding here is that whether ISPs really care about bad publicity or not, being blacklisted has direct effects on their operating costs as well as their quality of service. The latter may in fact drive customers away. As one industry insider described it: "A high cost action is to investigate each complaint rigorously. A different kind of high cost action is to do nothing."

Costs of infrastructure expansion

An incentive that was more difficult to gauge, is the effect of malware on the capital expenditures of the ISP – that is, the need to expand infrastructure and equipment as more spam or malware comes through

the network. A recent survey found that botnet-based denial of service attacks are growing faster in size than the ISPs are expanding their network – which is worrying the ISPs (Arbor Networks 2007).

Interestingly, infrastructure expenditures – apart from the costs of security equipment – were hardly identified during interviews as malware-related costs, a point to which we return shortly. As was mentioned earlier, interviewees pointed to customer contact as the highest security-related cost. When asked about infrastructure, a Chief Technology Officer answered: "The network is not affected. We have overcapacity to deal with DDoS. So that is not the problem." At another ISP, the Chief Information Security Officer told us: "We happen to have overcapacity of the network, so the growth in spam did not require us to expand the capacity." To which one of his colleagues added: "But the number of servers has increased, though." Others have argued that the volume of malware and spam-related traffic pales compared to the traffic from peer-to-peer networks and video streaming sites such as YouTube.com. We should add, however, that the presence of overcapacity may reflect the fact that we only interviewed ISPs in selected OECD countries. It may be different in other regions.

When we presented these findings to an expert in the economics of Internet traffic, he argued that our interviewees may be suffering from "the fallacy of the near." In his view, ISP employees dealing with security-related issues mention customer contact as their biggest cost because they are focused on the security budget, which includes the abuse desk as well as security-related customer support. To them the infrastructure cost "is just a number their accountant writes on a check every month." However, infrastructure is the main overall cost for any ISP, so any effect of malware on capital expenditures could potentially outstrip other expenditures. These costs do not gradually increase with the amount of malware and spam, but rather as a step function when capacity runs out. It is very difficult to relate these expenditures back to specific traffic patterns of spam and malware infections. Only higher up in the organisation are people in a position to compare the relevant numbers, although at that level the necessary security expertise and data is often missing. The interviewee argued that there are really three groups of people who all see a part of the problem, without being able to cross-connect it: "One group is dealing with malware, one group is dealing with the capital expenditures and engineering build-out and another group is dealing with handling the money." In terms of incentives, however, this lack of awareness implies that infrastructure cost is not a strong driver of the attempts of ISPs to reduce the impact of malware.

Benefits of maintaining reciprocity

An incentive that was mentioned by all interviewees is related to the informal networks of trusted security personnel across ISPs, CSIRTS and related organisations – which we mentioned earlier. When describing how their organisation responded to security incidents, interviewees would refer to personal contacts within this trust network that enabled them, for example, to get another ISP to quickly act on a case of abuse. There is not one informal network, but rather several overlapping ones. An ISP may approach a contact at a national CERT in another country so as to get in touch with the relevant person at an ISP in that country. These contacts are reciprocal. They are also contacted about abuse in their own network and are expected to act on that information. The incentive is that to maintain reciprocity, an ISP has to treat abuse complaints seriously, which is costly. The more abuse takes place on its network, the more other contacts in the network will ask for intervention.

Maintaining reciprocity not only establishes the informal network as a security resource, it also reduces the likelihood of being hit with blacklisting or other countermeasures. As one interviewee explained, "when we get in touch with service providers, we're saying, get this guy off the network or we're null routing your network from ours. If enough people do that, eventually they try to address security. The same thing happens if we have highly infected end-users hitting someone else, via malware or intentionally. What enforces security on a service provider is threats from other service providers." ISPs that are linked to the important informal networks typically get more leeway to deal with security issues

before significant blacklisting occurs. One ISP security officer told us that these informal contacts imply cost savings. Less staff time is needed to deal with the fallout of a security incident – *e.g.* going through time-consuming procedures to get off blacklists – and to deal with customer support.

Costs of security measures

So far we have discussed incentives that reinforce the benefits of security for ISPs with regard to malware. The incentive structure is mixed, however, and includes disincentives as well. An obvious disincentive is the costs of additional security measures. Typically, the tradeoff is between the direct costs of additional measures which are visible in the short term versus the costs generated by increasing security problems, such as customer support and abuse management. A security expert at a large ISP told us that for management it is difficult to estimate the amount of money the company may save with a technical solution which is supposed to reduce the costs of the abuse desk or call center. Another interviewee added that a complicating factor was that managers had encountered over-promising security providers who sold them 'magic boxes' that were supposed to solve everything.

We should mention, however, that the ISP's decisions often were not shaped by formal economic assessments or detailed analysis of their own cost structures. As one insider phrased it, "ISPs very much drive by the seat of their pants. Except for a very few of the largest ones, they are not actually examining the figures." When we asked how certain investments or measures were approved, the "business case" that supported them was typically rather commonsensical in nature, including rough estimates of direct costs and benefits, with the indirect ones not monetized or otherwise specified in any amount of detail. One interviewee told us that when considering security investments, they "look at the cost of *not* doing it" for which they produce rough estimates. Another ISP explained to us how they decided to set up a so-called 'walled garden experience' for infected users. Rather than disconnecting these users completely, the 'walled garden' provided them with access to security tools and Windows Update. A security officer explained the rationale behind this decision: "It costs a server or two. The rest of the stuff was free, we could reconfigure our existing infrastructure. The investments were the time that I put in. The 'walled garden' has a financial benefit because then the customer does not have to call as often."

Legal risks and constraints

Another disincentive is related to legal constraints. During the interviews, the European ISPs had different answers to the question on how much manoeuvering space the 'mere conduit' provision of the EU E-Commerce Directive allowed them. Monitoring their network more closely for security reasons could potentially lead to liability issues, some of the interviewees felt. In some EU countries, interviewees reported that privacy regulations that potentially treat IP addresses as private data had led their legal departments to set boundaries which affected the ability of the security staff to track malicious activity on their network – for example with regard to tracking individual IP addresses. One interviewee reported that security staff sometimes were not allowed to use information on malicious activity detected on the network. When asked about the limits of the 'mere conduit' provision, one security officer responded that they never encountered these limits, because the privacy regulations were much more constraining. Rather than monitoring their own network, this particular ISP could act on incoming abuse notifications for specific IP addresses and it relied heavily on this procedure. In a sense, the ISP was monitoring its own network through the incoming notifications from other ISPs, CSIRTs and the like.

Elsewhere there have been reports over liability issues around countermeasures, such as discarding the command and control traffic of a botnet or diverting it to where the botnet's behaviour can be studied more closely (Higgins 2007a). According to a security researcher "it involves mucking with a customer or peer's Internet address space... Obviously, liability in this area could be considerable." A security manager at a European ISP said "infiltrating is very risky and getting legal support for such matters, very difficult."

Some legal experts argued that these legal risks are non-existent, that they are based on an incorrect understanding of current legislation – *e.g.* that the EU data protection legislation does not at all conflict with network monitoring and other security measures. While that might be true, the reality is that the legal departments of some ISPs apparently interpret the situation – perhaps mistakenly – as rather ambiguous. These ISPs tend to be rather risk averse in dealing with this ambiguity. The transaction costs of clarifying these issues are, *ceteris paribus*, an obstacle to higher security.

Cost of customer acquisition

Other disincentives are closely related to the incentives discussed earlier. An interviewee at a large ISP mentioned brand damage as the reason why the business side of their company initially opposed the security measure to block port 25. They did not want to inconvenience their customers. Anything that might turn people away is a problem, because the cost of acquisition of new customers is high. The burden of proof fell on the security staff to convince management that the proposed measures were protecting the brand. Other ISPs also mentioned going to great lengths to avoid loosing customers while managing abuse. That might limit the effectiveness of their response to security incidents.

Rogue ISPs

Some of the security-enhancing incentives discussed above work as disincentives under different business models than those of the ISPs we interviewed. When dealing with abuse complaints becomes too costly, one can either reduce the amount of abuse on the network or one can reduce management of abuse – *i.e.* become less responsive to the complaints themselves. The same holds for customer support. In fact, such a lack of security could be part of the business model. It may, for example, allow an ISP to be cheaper than its competitors. One ISP indicated that a certain segment of its customers was actually "mini ISPs" which predominantly offered hosting services. The mini ISPs' retail prices were significantly lower than those of the upstream ISP from which they bought access, because they provided very limited support functions. Some of these mini ISPs would not patch their servers properly, thus becoming an easy target for malware. They were not very responsive to abuse complaints either. Our interviewee, being an upstream access provider, would then be contacted by other ISPs to take action against the mini ISP.

Another business model is sometimes referred to as "rogue ISP" or ISPs that are, in the words of one interviewee, "decidedly grey". These attract customers precisely because of their lax security policies. While these ISPs have more disincentives for improving security than the ones interviewed, they are not fully immune to some of the security-enhancing incentives we discussed earlier, most notably blacklisting. As one interviewee explained: "There are some ISPs in our country that are decidedly grey. They will take anyone and take no action against abuse. People will go there and then they will leave again, because they are unreachable [because of blacklisting]." Even rogue business models are eventually affected by blacklisting. "Suddenly, a Ukrainian ISP started answering our abuse reports," the interviewee continued. "Chances are that blacklisting had an effect on their business model. They are still not trustworthy, but it's a lot better."

An additional incentive for non-responsive ISPs is the pressure put on them by their upstream providers – the ISP "who feeds them the Internet," as one respondent phrased it. The higher up the stream, the more likely it is to find a provider who is in fact security conscious and sensitive to the incentives discussed earlier, such as maintaining reciprocity and blacklisting. In the example of the mini ISPs, their upstream provider forces them to deal with abuse complaints, because it reflects badly on the upstream provider if they do not. Beyond blacklisting, there is also de-peering – that is, an ISP may disconnect from a misbehaving ISP at an Internet exchange point. For the ISPs we interviewed, this is not an important incentive, because de-peering for security reasons is typically only employed against rogue ISPs, not

among regular ISPs. De-peering forces the disconnected ISP to buy transit service for its traffic, which implies much higher operating costs.

In sum

The balance between incentives and disincentives will vary depending on the ISP. On the whole, recent years have witnessed increased efforts by ISPs in dealing with malware, even in the absence of regulation or other forms of public oversight. The incentive mechanisms we discussed strengthen the ISP's own interest to internalise at least some security externalities originating from their customers as well as from other ISPs. In short, the current incentive structure seems to reward better security performance for legitimate market players – though it is sensible to keep in mind that in many countries price competition is intense, which is a disincentive with regards to security, other things being equal.

Several ISPs explained that they were at some stage of implementing technology that would automate the process of monitoring malicious behaviour on their network and quarantining the infected machines. One system monitored the network, cleaned malware from the traffic and automatically generated a list of 2 500 IP addresses a day of customers who have some form of security problem. When these cases hit a certain threshold, they would be automatically quarantined to only have access to Windows Updates and a range of security services.

While the technologies to automate the process of quarantining would help to scale up the ISPs response, it also brings into focus a critical bottleneck: the costs of customer support would become prohibitive if all infected machines were to be quarantined. A security officer at a large ISP estimated that the number of customers that would be affected at any time would be in the tens of thousands. While this number might go down over time as network security improves, it was obvious that the business side would not accept the cost impacts of such a measure.

Typically, the number of machines that are isolated on a daily basis is relatively modest – tens or, for large ISPs, perhaps hundreds of machines. At this level, the effort is effective in that it reduces the ISP's problems with abuse and blacklisting. But compared to estimates of the total number of infections on each network, these efforts look rather pale. When asked to assess the ratio between the actual number of infected machines on their network and the number of machines for which they receive abuse notifications, most interviewees estimate that the ratio is quite low. Only a small percentage of these machines would show up in abuse notifications and be dealt with. One interviewee called this "the two percent rule." A security expert was highly critical of the effectiveness of the efforts by ISPs: "Unless they are contacting more than 10 % of their customer base on a monthly basis, they are effectively taking no action".

A related issue is that the incentives of ISPs do not reflect the whole range of current malware threats. ISPs are predominantly sensitive to malware that manifests itself in ways that make their customers call in, leads to abuse notifications or that causes problems with blacklisting. That means spam proxies and DDoS attacks attract attention and raise costs, while spyware, for example, does not: "People get infected and it is very difficult to track them. Spam and DDoS is noticeable at the network level. But spyware stays on the computer, quietly collecting data." Others have argued that many ISPs are failing to prohibit the forging or spoofing of IP addresses by hosts as well as failing to filter outgoing traffic from IP addresses from which they are not authorized to originate.

Those security problems that are noticeable for the ISP will not always be addressed, either. Several ISPs mentioned "thresholds" of malware effects which needed to be crossed, before they would act on a customer's infected machine. Even then, the situation is often anything but straightforward. "The issue is, how do you help the people who are infected, given the current state of the security products in the market place? We see the traffic, we know there's something wrong, but how do you find what it is with the

current products? It's very hard... About 85-90% of the malware is not recognized by AV products, because a small change is enough to dodge the signature."

Another important caveat is that there are classes of ISPs for which the incentives to improve security are too weak or which even have strong disincentives to improve it, as discussed above. The ISPs we interviewed treat the existence of such ISPs as a fact. Because it is possible for rogue ISPs to stay outside the reach of legislation and law enforcement, they are going to be present for the foreseeable future. The ISPs we interviewed have learned to live with the presence of the rogue and semi-legitimate ISPs. They have found that they are able to operate quite effectively in this environment through a combination of tactics including those mentioned earlier, such as informal contacts that address upstream providers and blacklisting. In their mind, no matter what policies, governance structures or incentives are put in place there will always be providers, outside or inside their own jurisdiction, who will be a source of malware and other forms of abuse. Once this is accepted, then it is also accepted that an ISP has to build defenses and develop procedures for dealing with attacks. "You will always have to accept a certain level of noise, that is, of evil. You try to keep it below a certain threshold of irritation" said a security officer. This is one of the reasons why many ISPs are not impressed by proposals to regulate some set of baseline or best security practices for ISPs. One such proposal was under development by the Dutch electronic communications regulator OPTA but it was shelved for the time being after significant pushback regarding the legal basis for such regulations. The recent report of the UK House of Lords Science and Technology Committee (2007a, p. 31) also advocated making "good practice... the industry norm[, by means of regulation if necessary.]"

The fact that ISPs can work within the insecure status quo does not mean that their responses are static or complacent. The status quo actually contains significant incentives to improve security, which is why we have seen major changes over the past couple of years. Ironically, some of these changes, such as the policy to isolate infected machines, are not really advertised, for fear of dissuading customers from signing up.

E-commerce companies

The multitude of companies that buy and sell products or services over the Internet operate on a wide variety of business models, each with different incentive structures for security. We have chosen to focus particular attention on online financial services as they have been an important target of malware attacks, arguably more than any other sector (Counterpane & MessageLabs 2006). This includes brick-and-mortar banks who are offering part of their service portfolio online, credit card companies, as well as online-only financial service providers such as PayPal. The sector has been confronted with a wide range of threats ranging from botnet-assisted phishing spam runs and phishing websites to keyloggers and trojans that enabled man-in-the-middle attacks during secured banking sessions.

Benefits of increased online transaction volume

A key incentive for all these companies are the benefits they derive from processing online financial transactions resulting in a strong interest to enable a growing volume of online transactions. Credit card companies and online financial service providers typically charge a fee per transaction, either a flat amount or a percentage of the transaction. The situation is somewhat different for brick-and-mortar banks. For many of their services, they do not make any money from the transaction itself. Their incentive to pursue online banking is the considerable cost savings that it enables. Two of the interviewees in the financial sector estimated that online transactions were in the order of 100 times cheaper than processing those transactions offline, through their branch offices, mail or phone. Given the enormous volume of financial transactions, costs savings of that magnitude translate into a very powerful incentive to move online as much these services as possible.

How does this incentive affect security decisions? To answer that question, we need to understand how transaction volume interacts with several other incentives: the benefits of trust in the online services, the benefits of usability, the costs of security measures and the costs of fraud.

Benefits of trust

Within the sector, it is assumed that trust of consumers in the security of these services is a necessary condition for their uptake. That rewards investing in security. Beyond this generic consensus, however, views quickly diverge. There is disagreement over how strong the factor of trust is as a driver for the use of online services. Furthermore, it is unclear whether the current security problems with online financial services actually reduce that trust.

Several consumer surveys suggest that security problems turn people away from e-commerce and online banking in particular. The 2006 UK Get Safe Online survey reported that the fear of falling victim to Internet crime deters 24 % of respondents from Internet banking and has put off 17 % from Internet use all together (GetSafeOnline 2006). It is difficult to interpret the meaning of these findings when compared to other data. For example, most financial service providers still report significant growth rates in the adoption of their online services (PayPal 2007). These two seemingly contradictory pieces of evidence point out that the role and impact of trust is not yet adequately understood. An industry study of trust in e-commerce (Lacohée *et al.* 2006) argued that "[w]hile an initial hypothesis may be that people do not engage with online services because they do not trust them, our findings have shown that trust is not as significant a measure as first thought. What is more important to understand is that people are willing to take risks online, as long as they are informed, and it is clear how consequences will be addressed. People use specific services not because they trust them, but because they in some way provide a benefit to the individual and they know that if something goes wrong, restitution will be made." This suggests that an important factor driving the use of online financial services is not the level of trust in the security of these services, but the more specific expectation that a customer will be compensated in case of fraud. In other words, from a customer's perspective, it seems more important that financial service providers assume liability for online fraud than that they achieve a certain level of – perceived – security.

Benefits of usability

Assuming that increased security would increase consumer trust and, in turn, increase the uptake of online services this effect would still need to be weighed against the effects of increased security measures on the usability of the service. One of our interviewees at a bank with an international presence explained that the national branches of his company positioned themselves differently with regard to this tradeoff. While in some countries, two-factor authentication was readily accepted, in other countries the bank thought its customers were less open to such security-enhancing technology. If such measures were to significantly raise the threshold for people to do online banking, then the incentive to increase the volume of online transactions would influence decision making against such measures – even if this meant that fraud losses in those countries might be higher. By balancing usability and security, these companies try to maximize the growth of online financial transactions, while keeping the level of fraud at manageable levels.

Costs of fraud

Another important incentive for security is the fraud losses that accompany the increasing volume of online transactions. In the United States, banks are liable for direct fraud losses under the Electronic Funds Transfer Act of 1978 – also known as "Regulation E". Under this regime, customers are compensated for such losses, unless the bank can prove that the customer's claims are false. In many other jurisdictions, the banks are strictly speaking not liable for such losses. In practice, however, the banking sector has often

adopted voluntary codes which specify that customers who suffer losses are compensated – unless there are clear indications that they have colluded in the fraud.

To understand how the cost of fraud influences security decisions, it is important to look at some of the available numbers. The United Kingdom has arguably the best data available. APACS, the UK payments associations, publishes numbers based on actual banking data, not estimates based on samples and extrapolation. As one would expect, direct losses from phishing fraud in the United Kingdom have risen, though with a recent fall: from GBP 12.2 million in 2004 to GBP 33.5 million in 2006 to GBP 22.6 million in 2007 (APACS 2008). Over the past years the number of phishing attacks has increased significantly: from 2 369 attacks in 2006 Q1 to 10 235 in 2008 Q1. The broader fraud category of "card-not-present" fraud – which includes phone, Internet and mail order fraud – has risen from GBP 150.8 million in 2004 to GBP 290.5 million in 2007.

Not to downplay the seriousness of these losses, it is important to realize that the damage of phishing attacks is still well below the numbers for other fraud categories, such as stolen or lost cards (GBP 56.2 million in 2007) and counterfeit card fraud (GBP 144.3 million in 2007). Furthermore, while these numbers are going up in absolute terms, so is the number of customers banking online, as well as the overall volume of online transactions. APACS argues that the rise in card-not-present fraud should be viewed against the increase in the use of online or telephone transactions. While fraud has risen by 122 % from 2001 to 2006, the use of online or telephone shopping itself has grown by 358 %. Unfortunately, the available data is not sufficiently disaggregated to allow APACS to calculate fraud relative to volume. Credit card companies do publish such numbers. In 2006, VISA Europe reported that their overall fraud rate was at "an all time low" of 0.051% (fraud to sales by cards issued). However, card-not-present fraud, which includes online fraud, was the fastest growing type of fraud and now accounted for 40% of cases. PayPal recently reported their direct losses to fraud being 0.41% of overall transactions, but could not give information on the trend of their losses (House of Lords 2007b, p. 196).

Costs of security measures

While exact figures are hard to come by, the companies we interviewed all said their security investment levels are much higher than the direct yearly losses, often by one or two orders of magnitude. The capacity to deal with incidents is often already more expensive, let alone all of the preparatory measures and security defenses being put in place, such as the introduction of two-factor or three-factor authentication. The reason for this level of investment is that direct losses are not seen as representative of the overall problem. It would be much more devastating, for example, if online fraud eroded customer trust or slowed down the uptake of online financial services. Furthermore, there are reputation effects for banks that are targeted by attackers as well as for the industry as a whole. Nobody has robust estimates on either of these effects, which makes it difficult for financial companies to calibrate their security investments.

In general, the incentives are to keep fraud at acceptable levels and compensate victims, rather than to eliminate it. The latter would be economically inefficient, not only in terms of direct cost but more importantly because pushing fraud back further might require the introduction of security measures that make the use of online financial services less attractive to customers. A reduction in the growth of the online transaction volume is likely to imply higher costs for banks than the current damage caused by online fraud.

Companies, alone and through sector-wide collaboration, assess risks and prepare new security measures, which would rolled out when they feel the current defenses are no longer adequate. Exactly when is hard to specify. Some innovations have been put in place rather quickly. Phishing attacks, for example, are increasingly dealt with by contracting out response efforts to security providers who scan for phishing spam and hunt down sites that resemble the official bank website, at which time they initiate

notice and takedown procedures. Occasionally, this takes down legitimate web banking sites as well, when the security department is not aware of a marketing initiative from another part of the organisation and thus has not whitelisted the domain name.

Other innovations are deemed too early. In the Netherlands, there is an ongoing series of successful attacks on the two-factor authentication systems in place at most banks. Rather than introducing new structural new security measures, the banks have made incremental changes to the two-factor authentications systems which are relatively easy to defeat by the attackers. More structural measures, such as transaction authentication or three-factor identification, would require costly modifications to the back office systems as well as requiring their customers to learn new and more laborious security methods. So far, the response has been to make minor revisions to the existing systems so as to disable the last successful attack tactic. These measures are often accompanied by a number of other safeguards – such as temporarily slowing down the processing of real-time transactions. The direct financial losses of each attack have been relatively low, which makes the idea of the next successful attack less unpalatable. Ironically, one interviewee mentioned that the relatively modest losses per incident appear to be a deliberate strategy of the attackers. These attacks are trying to stay under the radar of the fraud detection systems – as well as making it less interesting for law enforcement to devote a large amount of resources to tracking down the criminals.

In sum

The incentives of financial service providers are such that in many cases they compensate customers for the damage they suffered from online fraud. They are willing to internalise these costs because the benefits far outweigh them. In that sense, they internalise the externalities of sub-optimal security investments and behaviours of their customers as well as the software vendors whose software is exploited to execute the attacks. Interviewees told us that when designing the security of their services, they have to assume that the end user PC is compromised. Many financial service providers claim they compensate all malware related losses. If that claim is accurate, then the security level achieved by the whole value net may not be too far from the optimum. The financial institutions bear the externalities, but they are also in a position to manage the risk through their security measures around online financial services.

However, there are several important considerations to take into account. First, one could argue that there are still externalities in the sense that important social efficiencies could be gained if people had higher trust in these services and would adopt them more quickly. These benefits could outweigh the additional security investments that would be needed. While the magnitude of these externalities is unknown, the financial service providers are the ones who stand to gain most from maintaining high trust in online services and, more to the point, from the increased adoption of these services. In other words, this is a problem of incomplete information, rather than of misaligned incentives.

A second consideration is that not all fraud-related costs to customers are compensated. While the financial institutions compensate victims for their direct losses, this might not cover all the losses that result from the fraud. In cases of identify theft, victims may not get all costs reimbursed and they may struggle for years with the consequences of having their personal information abused, such as blemished credit reports (TechWebNews 2005).

Third, in several countries the banking sector is reconsidering the existing liability regime, which might lead to liability dumping. Financial service providers have already started to push more liability onto the merchants. It seems we might see a similar trend for customers. Late in 2006, the Ombudsman for the German banking sector ruled against a customer who claimed to have been victimised by a Trojan, arguing that the customer provided no proof of a successful malware attack (A-i3 2006; Banktip 2006). The Ombudsman declared that the customer was not able to provide evidence of a successful malware attack

even though the customer's machine was infected with malware. This appears to shift the burden of proof onto the customer. In New Zealand, the banking association introduced a new code which has shifted at least part of the liability to customers. The new code allows the banks to request access to the customers' computer to verify that the operating system, the anti-virus software and firewall were all up to date. If this access is refused or the computer is deemed inadequately protected, the customer's claim may be turned down. Shortly after it was adopted the code drew severe criticism. In response, several banks and other stakeholders demanded changes that offer more protection to consumers. Currently, the debate seems to focus on the complicated question of determining just what part of the responsibility lies with consumers (South 2007).

The development of what one could call 're-externalising' fraud losses to the customers is not without risks to the banks themselves, as customer trust in Internet banking is partly based on the expectation that fraud losses are compensated. If customers experience more liability for their online transactions, it might reduce the uptake of these services, which directly affects the banks major incentive: the growth of online transaction volume. For this reason, a security official at a financial service provider called the attempts to shift part of the liability to customers "a very dangerous path to follow."

Ironically, the existing liability regime might actually be in the best interests of the bank. By internalising the damages, whether required by law or voluntarily, the banks have retained the freedom to balance the level of security against other factors, most notably the costs of security measures and the usability of online services. This allowed them to make more cost-effective tradeoffs than under a different liability regime. If they shift more liability towards their customers, then they run the risk of inviting more regulatory oversight for consumer protection.

One interviewee told us that while the US banks fiercely opposed the Electronic Funds Transfer Act of 1978 as it placed all liability on them, over time many in the industry realized that the regime was actually economically more rational for them. He called it "a blessing in disguise". Anderson (2007) found that during the period where the British banks operated under a more lenient liability regime for ATM withdrawals than the US banks, they actually spent more on security, as they were doing 'due diligence,' rather than actual risk reduction.

Some financial service providers argue that the current practice of compensating victims might provide a perverse incentive by rewarding customers for not securing their machine. Earlier experiences with ATM fraud suggest the risk of such a perverse incentive is manageable (Anderson 2007, pp. 13-16). Should banks re-externalise the costs of fraud to customers and merchants – or ignore potentially rising forms of damage which are currently not compensated, such as the costs of recovering from identity theft – then this might in the end lead to underinvestment or even overinvestment on the part of the banks when they invest on the basis of due diligence rather than actual risk reduction (Anderson 2007, pp. 13-16). In either case, the incentives would shift the level and type of security investments of the financial institutions away from the social optimum.

Software vendors

The very nature of malware focuses attention on software vendors. Malicious code exists because of software vulnerabilities that can be exploited – though we should not forget that there is also a class of malware that is based on social engineering, *i.e.* tricking users into voluntarily installing software that includes malware. The software market is highly differentiated although there are many linkages between segments such as operating systems and application software. Nonetheless, each market segment has somewhat different characteristics and hence creates different incentives for software vendors to improve security *ex ante* and *ex post* release and for malware writers to exploit vulnerabilities.

In recent years, much has been written about the incentives for software security. The predominant view seems to be that software markets do not reward security. In the words of Anderson and Moore (2007, p. 7): "In many markets, the attitude of 'ship it Tuesday and get it right by version 3' is perfectly rational behaviour." First, some authors claim that security is a "market for lemons", as consumers cannot tell secure from less secure software. One interviewee told us that he was in fact able to assess the security of the software his organisation bought, but that the different products were more or less the same in terms of security. So there was no real 'secure' alternative. Second, many segments of the software market tend to have dominant firms because of the combination of high fixed and low marginal costs, positive network externalities and customer lock-in because of interoperability and compatibility issues. "So winning market races is all important", Anderson and Moore conclude (2007, p. 7). "In such races, competitors must appeal to complementers, such as application developers, for whom security gets in the way; and security tends to be a lemons market anyway. So platform vendors start off with too little security, and such as they provide tends to be designed so that the compliance costs are dumped on the end users."

The analysis provides a powerful explanation for how we got to the current state of affairs. Its implications are less clear for what happens after the market race has been won by a software vendor. While any generalization is problematic, recent years have seen substantially increased efforts by many vendors to improve the security of their software. The development and deployment of vulnerability patches has improved. Arguably more important, the development of the software itself is increasingly focusing on security issues. Most of our interviewees agreed on this. They disagreed over the effectiveness of these efforts – some argued it was too little too late, others thought the market was moving in the right direction.

For obvious reasons, one cannot avoid mentioning Microsoft in this context. The company's problems and efforts have been most visible. By now, the story is well known. Given the market dominance of its Windows operating system, it has been a key target for malware writers. When the security problems plaguing the platform mushroomed early this decade, most notably in the form of global worm and virus outbreaks, Microsoft saw itself forced to change its approach. It all but halted development on its new operating system and re-tasked many developers to work on much-needed security improvements for its existing platform, Windows XP. These improvements were released in 2004 as Windows XP Service Pack 2 (SP2). While SP2 contained many vulnerability patches, it also introduced changes in the code base which set out to reduce the potential for vulnerabilities to be exploited. Furthermore, it turned on automatic updates and the Windows firewall by default.

For a variety of reasons, security among them, Microsoft then overhauled the code base for what would become Windows Vista, the successor to XP, at the cost of serious delays in the process. Vista's design introduced better security principles, which inevitably led to numerous compatibility problems when hardware vendors and independent software vendors had to adapt their drivers and programs to the new design. To a significant extent, the problems persisted even after the final release of Vista. Many would agree that these problems have slowed the adoption of Vista, as businesses and consumers wait for these problems to be resolved before switching. All of this implies substantial opportunity costs for Microsoft. There are no publicly available cost estimates, but it seems obvious that the security-related costs of SP2 and Vista are anything but trivial, even for a company of this size.

Microsoft is not alone in this trend reversal, though it might be the most dramatic example. In contrast, there are vendors who operate in markets that have demanded security from the start, such as the defense industry. These vendors have developed along a different path compared to those in the mass consumer market. As a result, their business models make it easier for them to economically justify security investments in the software development process. Just to be clear, the increased efforts in software security do not mean the problem of malware is getting smaller or even that the frequency with which vulnerabilities diminishes are discovered. There is a variety of factors at play, not least of which is end

users behaviour, which in combination determine if, how and when more secure software reduces the problem of malware.

Notwithstanding the different business models of software vendors, a number of incentives explain why this trend reversal took place. They point to the complex interplay between incentives and disincentives for security. Our findings do not conflict with the incentives mentioned in the literature. Rather, they confirm and complement them by focusing attention on incentives for security of established software vendors, *i.e.* after the "market race" has been won.

Costs of vulnerability patching

Developing patches for discovered vulnerabilities is costly, even if the fix itself is not hard to write. As one senior software security professional explained: "It's like the Mastercard commercial — two line code change, 20 minutes, finding every other related vulnerability of that type on every affected product version and all related modules, fixing it, testing it, 3 months. Giving the customers a patch they can use that does not break anything, priceless." Although it is daunting to calculate reliable and comprehensive numbers, the anecdotal evidence we were given suggests that an ongoing process of patch development, testing and release for a complex piece of software – like an operating system or an enterprise database system, which consists of tens of millions lines of code – is easily measured in millions of dollars.

Even more important, some interviewees argued, are the opportunity costs of tasking good software developers with vulnerability patching. One interviewee said: "If you reallocate the developer time for patches to other work, it might not be enough to build a completely new product, but you could build some complex functionality you could charge for. I could build something I could charge money for... if I did not have these defects to remediate."

Patching also imposes costs on the customer who applies the patch. This may include the cost of testing the patch before deploying it within the organisation, the actual deployment for all the relevant systems, as well as the costs of remediation when the patch turns out to "break something" – *e.g.* introduce system instabilities. Several studies have shown these costs to be substantial (*e.g.* August and Tunca 2006). Strictly speaking, the vendor does not experience these costs and some have suggested that these costs should be regarded as externalities that the vendor shifts onto its customers (*e.g.* Schneier 2007).

But there are indirect effects that do affect the vendor. First, it raises the maintenance costs of the software, which can be considered similar to raising its price and thus lowering demand – although this effect is significantly mitigated in case of lock-in effects or lack of alternatives. Many enterprises assess the so-called "total cost of ownership" of software, rather than just the price of the licence. It is not uncommon for maintenance costs to be much higher than the price of the licence itself. Second, if patching is too costly for customers, they may not keep their machines adequately patched. The resulting security problems may tarnish the reputation of the software itself – we return to brand damage and reputation effects shortly.

In response to these effects, many vendors have set out to reduce the costs of patching for their customers. For enterprises, patching is a different issue than for home users. The former need to have more control over the deployment of patches as patches potentially disrupt critical systems. In some cases, they might opt to not apply certain patches. "While it would be wonderful if everyone stayed fully updated all of the time," said one interviewee, "many enterprises choose to do extensive testing first, attempt to avoid blackout periods, and take into account many other considerations specific to their business before an update can be deployed. Enterprises that regularly deploy updates will be less vulnerable to malicious attacks, so with all of that in mind, each business must make the risk tradeoff that is appropriate for them."

The vendors we spoke to described efforts to better support their business costumers in this regard. Microsoft, for example, introduced Windows Server Update Services (WSUS) which allows IT administrators to control the deployment of patches across the computers in their network. Furthermore, vendors try to improve the information they provide with patches, so that business can make an informed risk assessment regarding if, when and how to deploy a patch. Several interviewees also indicated that enterprise customers asked for bundled patches, which are tested and released together on a regular schedule (*e.g.* weekly, monthly or quarterly), rather than single-issue fixes that are released as soon as they are ready. "We do not do single fix patches, it's not economical and you cannot keep the quality up", said one interviewee, adding that some of their customers even wanted the frequency of patch releases reduced to twice a year, so as to decrease the costs on their end.

For home users, reducing the costs of patching has mainly consisted of developing easier, more user-friendly mechanisms to deliver and install patches. Microsoft developed "Automatic Updates" and turned it on by default in XP SP2. The vendor reported that over 350 million Windows machines world wide receive the monthly "Malicious Software Removal Tool" through Automatic Updates or Windows Updates (Microsoft 2007). In the environment of open source software, Firefox – an Internet browser with the second-largest market share, after Microsoft's Internet Explorer – has enabled automatic updates by default since version 1.5. Rather than bundling patches, the developers of Firefox release the patches as soon as they are ready. The default setting of the browser is to download and install them at the earliest opportunity. The developers recently reported that under this new model, 90 % of Firefox users installed a recent security patch within 6 days (Snyder 2007).

The costs of patching could also work as a disincentive to security for those vendors that seek to avoid these costs. As a result vulnerabilities remain unpatched for too long, if they ever get patched, or the quality of patches might be too low. The urgency of this issue increases if attackers are indeed, as has been reported, moving way from exploits in the operating system and toward third-party applications and hardware drivers (Lemos 2006). However, not providing vulnerability patches does not seem to be a tenable strategy for an established vendor whose product is actively being targeted by malware writers. On the other hand, even substantial efforts in patch development can leave a software product vulnerable – *e.g.* because patches are more complicated to develop and test for products that are tightly integrated in a larger software package. An analysis of the known vulnerabilities for Internet Explorer found that for a total 284 days in 2006, there was exploitable code available for known, unpatched critical flaws in Internet Explorer 6 and earlier versions (Krebs 2007).

If a vendor's market position requires to performing costly patch development, then these costs might incentivize more investments in security early in the development process, in the hope of reducing the number of vulnerabilities after release – or perhaps more accurately, the frequency with which these are discovered.

Costs of secure software development

While vulnerability patching is generally seen as desirable, even if not by everyone (Rescorla 2004), many have argued that it does not really solve the underlying problem. Finding and patching vulnerabilities might not make the software product itself more secure. Some research suggests that for many products, the discovery rate of bugs is more or less constant over time – in other words, finding and fixing a vulnerability does not reduce the chance of an attacker finding a new vulnerability to exploit (Rescorla 2004). Furthermore, patch development consumes resources that are not used to make software more secure before it is released.

This is a valid criticism. However, several interviewees made the case that costly patching procedures still provide an incentive for more up-front investments in secure software development. One argued that

the more powerful incentive for secure software development is the fact that back-end patching costs are much higher than the costs of preventing the vulnerability during development. Another interviewee told us: "The argument to make for writing better code is cost avoidance, even if you charge for support (and we do). The way you get a good margin on it is if you can charge for maintenance but you do not have to constantly produce patches because those are expensive, that cuts into your margin."

We did not come across economic analyses that directly compare the costs of secure development with those of patching. It is unclear whether vendors even have this kind of data available. One interviewee told us: "I cannot add up what we've spent on the front-end... Most of secure development is good development, not some special security add-on." It seems clear, however, that the costs of secure software development are substantial. It requires more resources and can affect time-to-market of a new product – a critical factor in many software markets though here too the effect may be tempered by customer lock-in. Furthermore, secure development often involves costly assurance processes. One interviewee described the so-called "Common Criteria" evaluations for major releases of their products. These evaluations are made by external consultants and were estimated to cost between USD 0.5-1.0 million each – not including the time-consuming involvement of internal staff.

Even in the absence of hard numbers, the interviewees were adamant that there are significant cost savings to be made by investing in secure software development. After Microsoft started its "Security Development Lifecycle" initiative, it published some preliminary numbers which appeared to support the idea that the new approach resulted in significant reductions in the number of vulnerabilities found after release (Microsoft 2005). In addition to reducing the direct costs of patching, there are reductions in opportunity costs which potentially are even higher. In the words of one interviewee: "I worry about the opportunity cost of taking good developers and putting them on tasks for security patches for avoidable, preventable defects. That's why we put a lot of work up-front to avoid that. We have training, we have automated tools – anything you can do earlier in the cycle is goodness. It's never been hard to justify those costs."

Cost of brand damage and reputation effects

An additional explanation for the increased security efforts of software vendors are the reputation effects that they suffer for poor security – or enjoy for good security. The strength of these effects are notoriously difficult to estimate. Some have suggested that they provide a fairly weak incentive (Schneier 2007). Whether that is true or not, it does seem to play a role. The major security-related changes within Microsoft were driven by the major worm and virus outbreaks in 2002 and 2003. The key difference between those security incidents and ones that preceded them was scale and the resulting damage. Neither affected Microsoft directly. The reputation effect of those incidents seems to be the most plausible explanation for the changes in the company's course.

As mentioned earlier, Microsoft has invested in mechanisms to make it easier for its customers to patch their machines, even though they do not suffer the customer's patching costs directly. Furthermore, so far Microsoft has allowed pirated versions of Windows to download security patches. This appears to value the reputation of the platform higher than denying services to non-customers. Keeping their customers patched as much as possible helps to reduce the scale of security problems that the platform is associated with.

The incentive of reputation effects might be stronger in open source communities, where reputation is a very valuable resource (*e.g.* Watson 2005). It might help to understand why early in the development of what would become the Firefox browser – shortly after the code of Netscape Communicator had been open-sourced in 1998 – the developers made a number of security-conscious choices. The security

performance of the browser played a key part in its marketing the positive evaluations of software reviewers.

While there are indeed incentives that help us to understand the intensified efforts toward security, they are also counteracting incentives, which complicate the drive towards more secure software. These incentives help us to understand why despite increased efforts, making software more secure is difficult under current market conditions.

Benefits of functionality

"Part of the reason for the mess is that people want fancy gadgets and do not care as much about security, and that's exactly what they got," one software security professional told us. The 'gadgets' referred to in this statement are the functionalities provided by software products. Even vendors with an established market position will at some point want customers to buy a newer version of their product or a complementary product. Another interviewee said: "No-one buys your product only because it is secure, they buy it because it allows them to do new things." The drive of the market to produce ever more powerful software has generated numerous innovations. At the same time, it has made it much harder to build secure software.

Functionality versus security is not necessarily a zero-sum tradeoff. New functionality can be security related, for example, or it might be implemented securely. In practice, however, they can be difficult to reconcile. The history of software development is rife with examples where tradeoffs in the design of software have often favoured functionality over security. Many of the much-maligned features of Microsoft's Internet Explorer, such as its deep integration into the Windows platform, started out as functionality – *e.g.* the ability of a website to silently install code on a user's system, which would increase the functionality of the system without requiring the user to understand and manage the process of installing software. There have been many beneficial uses of this functionality, but it also has turned out to be a huge security risk. In response, IE7, the latest version of Internet Explorer, has reversed many of these design decisions.

There is an intrinsic tension between adding functionality and making software more secure. Security benefits from simplicity and a limited amount of code (*e.g.* Barnum and Gegick 2005; Bernstein 2007). Many of today's major software products are neither. The need to expand functionality with each release only exacerbates the situation. Of course, secure software development practices set out to mitigate this problem, by reducing the "attack surface" of a certain functionality and manage the remaining risks or, if the functionality is inherently insecure, to exclude it from the product.

One could argue that as the security-related costs of users go up, the market will reward security-related functionality that can reduce those costs. There are several well-known counter-arguments to this – including lock-in effects, lack of alternatives, weak market signals for security and the information asymmetry between vendor and customer. That said, there appears to be a market demand for certain security improvements, most notably those that reduce the total cost of ownership. Some software products, both proprietary and open source, are actively marketed as being more secure and less costly to maintain than their alternatives or predecessors. Whether the market over time can distinguish between empty claims and security improvements that actually achieve cost-savings is not yet clear.

Benefits of compatibility

As discussed above, software products benefit from positive network externalities. The value of a software platform – such as an operating system – increases non-linearly with the number of users. There are two sides to this: the more users there are, the more vendors will want to develop software for that

platform; and the more software there is for the platform, the more users will want to adopt it. Anderson and Moore (Anderson and Moore 2007, p. 5) concluded that all of this implies that platform vendors will impose few security restrictions so as to appeal to third party software vendors – *i.e.* to maintain compatibility and inter-operability of software. How these incentives play out of for a specific vendor depends on the type of product they provide and the position they have in the market.

For a dominant platform, maintaining compatibility is key when moving from one version to the next. As one industry insider told us: "The only thing [Microsoft] cared about in the transition from Windows 95 or Windows 98 to Windows XP was application compatibility, otherwise people would never move to XP." This had all kinds of effects on security and the problem of malware. To achieve maximum compatibility, the default installation of XP set every user up with administrator privileges, which means that people typically operated their machine under a user account that allowed unrestricted control over the machine. From a security standpoint, this is undesirable, because it means that once a machine is successfully attacked during use, malware has full access to the machine and can, for example, apply changes to the operating system and install root kits that are incredibly difficult to detect and clean up. Better security practice would be to set up an administrator account to be used only when new software needs to be installed or system changes need to be made. The rest of the time, users should run by default as standard users, with restricted privileges. This reduces the "attack surface" – *i.e.* the amount of code, interfaces, services, and protocols available to an attacker.

In response to the default user setup of XP, third-party vendors assumed that all users would run with administrator privileges and they designed their programs accordingly. In turn, because so much software assumed the user ran with administrator privileges, running the system as a regular user with limited privileges was not really viable. "The end user was pretty much forced to run as administrator", said one interviewee. While they might not have much of a choice, end users were accustomed to having full control over their machine, unbothered by security restrictions.

Large organisations did sometimes set up the desktops of their employees with restricted regular user accounts. This was is a costly set up, however, because it requires a lot of support staff to manage these installations. Even minor changes needed administrator privileges and thus a support staff action. Of course, if you set up your users as administrators, the support costs are also high, because of the increased security risks.

The only way to break out of this self-reinforcing path dependency is for everyone to adapt their behaviour. During the development of Vista, Microsoft decided to change the default way user accounts were set up. This required Microsoft developers to create a viable standard user mode with restricted privileges. They introduced User Account Control (UAC) for this purpose. Their enterprise customers, many of whom wanted to run their desktops under standard user accounts, applauded this development, as it promised to reduce their total cost of ownership. The problem was that it created serious compatibility issues with the existing third-party software, much of which still presumed administrator privileges. While vendors were informed about the upcoming changes, many did not actually adapt their code to work with these features. One interviewee explained that it was not attractive for vendors to comply with the new restrictions, because they had to invest in changing their code just to get the same functionality that they already had before Vista.

When Vista was released, a substantial number of these compatibility issues were unresolved, even though Microsoft itself developed auto-mitigation measures to deal with many application compatibility problems that the vendors did not resolve themselves. Users experienced poor or missing device drivers and incompatible software programs. Many complained about the constant security prompts and warnings that UAC confronted them with. Because many programs did not run properly in standard user mode, they constantly had to ask for elevated privileges, which triggered the UAC prompts. This was exacerbated by

the fact that UAC was not implemented very elegantly and thus generated more prompts than needed. As one interviewee explained, the move to UAC "is considered a paradigm shift that can translate into worse user experience if the user is running software that has to elevate every day."

Microsoft anticipated these problems to a certain extent. They felt that the compatibility problems of end users were worth the price in moving the software industry to build software that could operate under a standard user model. But without a way to force the third-party vendors to adapt their software, this would be "a dangerous game to play," said one interviewee, as Microsoft itself will receive part of the blame for these problems. UAC is one example. Other security improvements in Vista suffer from the same incentive problem: They only work if the independent software vendors adapt their code. If using the security feature is not turned on by default, the vendors might simply ignore it, which means that the feature does not actually improve security for end users. If the feature is turned on by default or if it cannot be turned off, then users will experience serious compatibility issues. These compatibility issues likely translate into a postponed adoption of Vista, especially by enterprise customers, as they wait for these problems to be sorted out before they move to the new platform. For Microsoft, postponed adoption means that pushing the market towards these security improvements imposes substantial opportunity costs.

On the whole, the benefits of compatibility and inter-operability create strong path dependencies which can only be broken away from at high cost.

Benefits of user discretion

An issue that runs throughout the challenge of software security is user discretion – that is, key decisions about how to configure and operate the software product are left to the user. The user – or in enterprise contexts, the IT administrator – decides whether or not to install vulnerability patches, the user decides whether to operate within User Account Control or to turn it off, the user decides how to configure a firewall, and so on.

User discretion allows software products to be adapted to a wide variety of contexts and user preferences. That means the product can reach a wider market and can create more benefits for its users, making it more valuable. Perhaps more importantly, user discretion touches on property rights. Software runs on machines that are not owned by the vendor. In principle, it's the owners who should be able to decide how to balance tradeoffs between functionality, performance, availability and, yes, security – as well as any other value relevant to them. After all, the owners are the first to bear liability for what their system does – whether this affects themselves when patch deployment breaks critical business applications, for example, or others, when their systems are compromised and used to attack other users. "We are not in the business of telling our users what to do," was how one interviewee summarised it. "We can inform them, educate them and provide them with the appropriate tools, but we cannot make these decisions for them."

With user discretion comes user responsibility. This is a blessing and a curse for software vendors. The blessing is obvious: many of the current security problems fall within the realm of user behaviour rather than within the realm of software production. This shields vendors from part of the responsibility to resolve these problems. Of course, it is also a curse. The decisions that users make affect the security performance of a product, which in turn affect the reputation of the product and its vendor. There is plenty of evidence demonstrating that in many cases, users lack the information or expertise needed to make rational security tradeoffs or that their decisions do not account for the costs they impose on others – including, but not limited to, reputation damage to the software vendor.

There are limits to user discretion. There are hard limits, where software simply does not enable or allow you to take certain actions, and softer limits, where the default configuration of a product tries to

guide behaviour in a certain direction. For example, when Microsoft introduced UAC, it turned the feature on by default, but it did include the possibility to turn it off by changing the system settings. Preliminary feedback indicates that, so far, over three quarters of users keep UAC turned on.

Where and how to set such limits is a difficult balancing act for vendors. It implies many tradeoffs between user discretion and protecting the integrity and reputation of the product. As one interviewee explained: "That debate raged on for four years straight, from the team level to the senior VP level and we rehashed that debate fifty times in those four years. You know – what should the defaults be and how much pain can we put the users in to get through to the independent software vendors? Are we being too aggressive with this plan or are we not aggressive enough? It was a huge engineering decision that really took a lot of guts at the VP level to support because we knew we were going to generate some customer dissatisfaction. But the alternative is to say: I hope anti-malware engines can keep up with malware."

Sum

Software vendors work under a mixed set of incentives which may vary for different market segments. They do experience increasing costs as a result of growing security problems, most notably the direct and indirect costs of patch development and reputation effects. That explains why many vendors have substantially increased efforts to improve the security of their software. The vendors also experience incentives that make it costly and difficult to introduce more secure software, even if they are willing to invest in development. The net effect of the mixed set of incentives is dependent on the product and the market segment in which the vendor operates. Assuming all other things are equal, the increased efforts mitigate software-related security problems. However, at the same time as security efforts are being increased, malware is becoming more sophisticated, adapting to the new defenses. Notwithstanding the efforts of software vendors, many of our interviewees expected that the situation would get worse still, before it would get better.

Vendors do not bear the full costs of software insecurity – *i.e.* there are externalities. Schneier (2007) has repeatedly argued that all the money that consumers of software products are spending on additional security products and services should be counted as externalities generated by those software products. That might not be fully correct and overestimate the size of the problem. To a certain extent, security problems are connected to users' decisions and behaviours – as is inevitable, given user discretion over the configuration and use of software, as well as social engineering attacks which do not need software vulnerabilities to compromise a system. If somebody decides to buy a cheap or highly functional software product with known security problems plus separate security software, it is that consumer's choice and should not be treated as an externality. In theory, a well-functioning market would offer software with different degrees of protection and let consumers choose. However, that assumes that everybody has full information and that there are no externalities on the consumer side. As we know, in many software markets consumers experience lock-in effects or a lack of alternatives. So there are externalities generated by the vendors' decisions, but they are probably lower than the total cost of security measures.

Registrars

The Domain Names System is part of the Internet infrastructure and as such it is affected by malware in a variety of ways. There have been highly publicised botnet-assisted DDoS attacks on root servers and TLD name server operators, aided by sophisticated tactics that employ the existing DNS infrastructure to amplify the attacks.

In addition to these threats to the DNS infrastructure posed by malware, new attacks which combine phishing with compromised web servers or end user machines – such as so-called 'rock-phish' attacks and 'fast-flux phishing domains' – have pulled the registrars more directly into the fight against malware. The

fight against phishing is led predominantly by market players who are targeted by the attacks – *i.e.* banks, e-commerce companies, etc. – or by security service providers working on their behalf, often assisted by expert volunteers working at ISPs, CSIRTs and other organisations.

The procedures to take down phishing sites are changing constantly, as attackers adapt their strategy in response. Typically, ISPs and registrars are involved in taking down a phishing site. The first takes down the hosting website, while the latter removes, suspends or redirects the domain names used by the attackers. Redirecting a domain name means sending the traffic to another location, typically to allow law enforcement or security specialists to examine it more closely. Suspending is sometimes preferred over removing it, as the latter would allow the attacker to register the name again elsewhere. The response of ISPs and registrars to the notification of phishing sites varies. Some act swiftly, others do not. At the latter extreme, we find bullet-proof hosting, whose business model is based on non-response and keeping malicious sites online as long as possible. Research suggests that legitimate ISPs and registrars, once they are under pressure to act, go through a learning process and develop procedures to deal more swiftly with abuse (Clayton 2007). At that point, the criminal activity starts to migrate to other, easier targets.

The transaction costs of domain name registration itself are very low – as evidenced by the practice of "domain tasting," where millions of domain names are registered, the overwhelming majority of which are cancelled before the so-called "grace period" expires. For the registrar, this process is profitable because it enables a business model to find profitable domain names through trial and error, which drives up the number of registrations that do make it past the grace period and thus generate revenue. Some interviewees suggested that there is a relation between domain tasting and malware, but within the context of this study we have been unable to find sources to clarify and corroborate that relation.

The incentives of ISPs have been discussed earlier. What about the registrars? To a significant extent, ISPs and registrars are overlapping categories. Domain name registration is an extremely low margin business, which is why many registrars tie them to complementary conventional ISP-type services, such as web hosting and hosted e-mail services. Some registrars even offer domain names at a slight loss, in order to entice people to register through them, knowing that a portion of them will sign up for complementary services. For the registrars that do not offer complementary services, it becomes a bulk business in order to survive solely on the very small margins of domain name registration.

The overlap between registrars and ISPs means they share similar incentives. It also means that the size of their operations is such that staffing an abuse desk and other security-related positions is seen as a normal cost of doing business. The different parts of the business often share a centralised abuse desk. Furthermore, they need such capabilities for other reasons than just security, most notably to deal with complaints regarding copyright infringement – our interviewees reported that the latter made up a large portion of the incoming complaints. Of course, there are also smaller registrars, with or without complimentary services, who lack staff to deal with abuse – again, similar to the situation with ISPs. Some of these smaller registrars leave it to the hosting provider to deal with all content-related complaints. Because of the overlap between registrars and ISPs, we refer back to the section on ISPs to get a sense of the incentives that both have in common. We only briefly summarise them here, complementing them with more specific findings for registrars.

Costs of customer support and abuse management

As with any business in a competitive market, registrars have an incentive to reduce operating costs. This includes customer support and abuse management. The number of complaints was reported to have risen substantially in recent years, though part of this growth coincided with growth of the customer base. At the same time, the response process has become partially automated and thereby more efficient. To illustrate: one interviewee reported getting 1 200-1 500 incoming complaints per day for a customer base

of several million. Only a minor part of the overall incoming notifications relate to malware. The bulk consisted of complaints about spam or copyright violations. While the company in question offered complimentary services, most of the incoming complaints were about domain names that were registered through them, but hosted elsewhere. They were contacted because their terms of service did not allow the domain to be used for any kind of abuse – and they have a reputation for enforcing these terms. On the whole, the interviewee estimated that they suspend around 20 domain names per day for abuse-related reasons. Only a few per week were specifically for malware. One explanation offered for this relatively modest number was that for end users who were infected by malware, it is often difficult to tie that infection to visiting a specific hosted domain.

With the core process of registrars being relatively low cost, involvement in notice and takedown procedures can drive up operating costs. Dealing with abuse notifications requires staff. The cost of collaboration therefore provides an incentive that, ceteris paribus, works against security. This is reinforced by the need to investigate the notification, to understand whether the domain name is indeed associated with malicious activity. Given the dynamic and increasingly sophisticated strategies of phishing gangs, this can be more difficult than it may seem at first glance. Even for the experienced staff at larger registrars, investigating a notification and request to suspend a domain name for malware related issues can take several hours. Phishing sites are less difficult to investigate and can typically be dealt with within an hour.

The incentives for criminals are to register with registrars who are slow to respond to abuse. The longer the domain name stays active, the more successful their attack can be. This means that not all registrars are equally affected. Those that are swift to suspend, remove or redirect a domain name typically incentivize criminals to look for easier targets. Given the enormous variety of registrars, both for generic and country-code top-level domains, an easier target is usually not hard to find. These registrars do experience consequences for their lack of responsiveness, similarly to the consequences that ISPs suffer. In that sense, the costs of customer support and abuse management work as an incentive to improve security. Our interviewees explained that it was their experience that if they dealt proactively with abuse, then criminals would avoid them or move elsewhere, which reduced the amount of complaints coming in, as well as associated costs such as blacklisting. The amount of abuse had gone down relative to the growth in their customer base.

Costs of blacklisting

In as far as the registrars offer hosting and e-mail services, they are subject to the issue of blacklisting along the same lines as the ISPs. Blacklist operators also watch registrars and their responsiveness to abuse complaints. In extreme cases, blacklists may be directed at the registrar itself. A case in point is the recent row between the blacklist operator Spamhaus and the Austrian registry/registrar Nic.at. Spamhaus had requested Nic.at to remove several domain names it said were associated with phishing by the "rock phish" gang. Nic.at did not comply with these requests, citing legal constraints. They argued that they could not legally remove the sites, unless Spamhaus provided them with clear proof that the domain names had been registered using false information (Sokolov 2007). The conflict escalated when Spamhaus added the outbound mailserver of Nic.at to one of its blacklists – listing them as "spam support" – so that the registrar's e-mail was no longer accepted by the multitude of servers using this popular blacklist. About ten days later they changed the listing of Nic.at to a symbolic listing – no longer actually blocking the IP addresses, but keeping them listed as "spam support." Several of the offending domains have been removed, but Nic.at denies that they complied with the request and assumes that the hosting providers took action (ORF 2007; Spamhaus 2007).

Benefits of maintaining reciprocity

For registrars, maintaining reciprocity is as important as it is for ISPs. We heard numerous examples where registrars with hosting and e-mail services could prevent instances of blacklisting through informal contacts with blacklist operators as Spamhaus as well as major e-mail and network providers. One interviewee mentioned that one direct benefit of being responsive to abuse complaints is that it typically keeps sites with security problems off blacklists – or at least ensures a proportionate response from blacklists, such as listing the specific machine associated with the abuse, rather than listing a wider range or subnet in which the offending machine resides. A security expert at an ISP claimed that his organisation sponsored Spamhaus, which effectively gave them a free pass in terms of being blacklisted.

An interesting example of reciprocity that was added to earlier findings was the issue of size of the customer base. According to one interviewee, the larger the hosting provider, the less likely it was to get blacklisted by the large e-mail providers such as AOL, as it affects AOL's customers as well when they cannot reach websites or mailboxes at the hosting provider. This effect is far less likely with smaller connectivity, hosting and e-mail providers.

Legal risks and constraints

As with the ISPs, a number of legal ambiguities surfaced which in some cases translated into disincentives for security. Some interviewees argued they had to be careful with monitoring the hosted sites on their network. One interviewee said: "The legal liabilities kick in as soon as you have knowledge or should have knowledge that something took place on your network. If you are proactively monitoring all the content of your hosting customers but for whatever reason something is missed, while there is an expectation that you should have caught it, then you could potentially be held liable for that content. So the monitoring that we do is somewhat limited in scope and only applies to areas where there is some sort of a safe harbor legal provision."

Then there are potential liabilities around suspending or removing domain names, as it involves a contractual relation between registrar and registrant. Even if the terms of service of the registrar preclude the domain name being used in relation to spam or other forms of abuse, that still requires the registrar to investigate and build a case showing that those terms have been breached. That can be costly. Several interviewees in the security community pointed out that security professionals often use a short cut: rather than asking the registrar to adequately investigate and decide on an abuse complaint, they point out that the registrants WHOIS information is false. As one interviewee explained: "For those registrars that are not willing to assume the risk of the liabilities, the WHOIS accuracy policy is a comfortable refuge." Referring back to the case of Spamhaus vs. Nic.at, the request of Spamhaus was indeed to suspend the phishing domains on the grounds that their WHOIS information was false. The response of Nic.at was that they were contractually bound and unable to remove the domain names unless Spamhaus could provide legally meaningful evidence that the WHOIS information was indeed false.

There is also the risk of collateral damage from removing domain names. It could be that the domain name is indeed used for phishing, but that not all activity associated with it is criminal or that the actual owner is unaware of what is going on. The fact that the registrar acted in good faith upon the request of others would in all likelihood not shield it from liability, unless the request had a legal basis, such as a formal request from a law enforcement agency – ignoring for the moment the obvious complications that would arise should different national jurisdictions be involved. Early 2007, registrar GoDaddy.com received a lot of criticism after it removed the DNS record for the security website SecLists.org at the request of MySpace.com, after the security site published a list of 56 000 MySpace usernames and passwords that had been circulating on the Internet (Utter 2007).

Even if the domain is actually owned by criminals, that does not mean the registrar is shielded from repercussions. In the past, there have been cases of spammers successfully suing their ISPs for shutting them down, just as they have sued blacklist operators such as Spamhaus – a case which was initially won by the spammer, although that did not affect Spamhaus directly because it is located outside the courts' jurisdiction. In short, the risk of liability drives up the costs of compliance with abuse notifications, especially in combination with more complicated and difficult to diagnose attack strategies, which work against security.

Not everyone agreed that these liabilities form a significant risk. "In a lot of cases the risk of incurring liability vis-à-vis a spammer or malware author is very minimal," said one interviewee. "I believe most registrars operate on that premise. Certainly, I have heard the excuse of liability used by some registrars and I feel that it should not be used to absolve yourself from your responsibility to your customers and your community... The real risk is the cost of defending yourself against court cases. Even in the most ludicrous cases there is some exposure and you need to take those exposures into account into your business model."

Costs of brand damage and reputation effects

On the positive side, there appear to be reputation effects which provide security-enhancing incentives. As mentioned earlier, there are several cases of registrars who were popular among phishers and who at first did not respond to requests to suspend domains. Then they apparently went through a learning process and started to remove domain names quickly in response to requests (Clayton 2007). It is unclear what precisely prompted this learning process, but their behaviour suggests that the registrar does not want to be associated with the malicious activity. Another case is the ccTLD of Tokelau, an island with 1 300 inhabitants associated with New Zealand. The registrar for the .tk domain is a Dutch-American company, which hands out most domain names for free, making money from showing advertisements on the registered domains. After McAfee announced that over 10% of the .tk domains were suspected of malicious activity, the registrar introduced new measures, which included frequent scanning of the domains for malware (Dot-TK 2007).

Costs of customer acquisition

Interviewees expressed mixed views about the effects of the costs on security of acquiring and retaining customers. The dominant view appeared to be that proactively fighting abuse actually helped to acquire and retain customers, as it helps build their brand as trustworthy and secure. In addition, active abuse management helped the registrars to mitigate risks of blacklisting, also for customers that were not directly involved in the abuse issue. Non-responsive registrars and hosting providers might experience more severe forms of blacklisting which are correlated with substantial collateral damage within their customer base.

The other side of that story is that proactive abuse management often implies swift action, which might be perceived as hasty or unjustified by the customers involved in the abuse issue. The latter might see themselves as victims of the abuse management as well as of the actual abuse. In general, the organisations we spoke to take great pains to resolve abuse situations without alienating the customers – with the obvious exception of those customers who are in some manner complicit.

Sum

Registrars face a mixed incentive structure for security that varies across the different business models. To the degree that registrars operate as ISPs – and many do as they tie in registration services with hosting e-mail and other complementary services – they face a similar incentive structure. There is some

evidence that suggests that registrars are indeed responsive to outside pressure and that improved security provides benefits (*e.g.* Clayton 2007). A security officer at an international bank told us he was not worried about the fast-flux networks for phishing, because in his experience registrars were quite responsive in addressing the attacks at the level of the domain name. That still implies, however, that in the absence of outside pressure, the incentives for security are not strong. In light of the large number of registrars currently in operation, this suggests a long learning process, even if we assume that registrars which have improved security will not fall back into complacency.

As was discussed earlier, the abuse complaints that ISPs receive cover only a fraction of the actual amount of abuse on their network. The interviewees confirmed that this is similar for the domain names or hosting services that fall under their purview. "For every abuse situation we are notified about, there are probably several more going on that we do not get notified about," said one interviewee. In practice, this means that while many registrars may have incentives to improve security, their efforts do not reflect the full extent of the security problems associated with their services and their customers. In other words, there are externalities arising from these services for other market players in the value net.

End users

End users are arguably the most heterogeneous set of market actors, ranging from average home users to SMEs to public institutions to global corporations. Rather than trying to differentiate all of these actors, we briefly discuss two extreme categories – home users and large organisations, public and private – and discuss in general terms the incentive structures under which they operate.

Home users

The rise of botnets has turned the problematic security practices of home users into a collective problem. Home user security has never been strong, but until a few years ago the consequences of this behaviour mainly affected the users themselves. That incentive structure has changed dramatically. By masking its presence to the end user, malware can turn end user machines into attack platforms which are used against many other players in the value network.

The lack of end user action against the infection of their machine is a combination of incomplete information – not knowing that they are infected or unable to evaluate the relevant security risks and defense strategies – and incentives – not bearing the costs of their decisions to others. Incomplete information is important, because it further weakens the already misaligned incentive structure. While it is true an infected machine is often mobilised for use against other actors than the machine's owner, it is certainly also true that a significant portion of malware poses a direct threat to the owner – for example, keyloggers that capture access codes to financial accounts, 'ransomware' that renders user files inaccessible until a ransom is paid to the criminal or Trojans that enable man-in-the-middle attacks during secured online banking sessions. In principle, these risks could provide a strong incentive for home users to secure their machines. But their lack of understanding of such risks or how to defend against them renders the incentive to act on them rather weak, if not inexistant. The interviewees at ISPs told us that when they contact users whose machines have been compromised, the response is generally quite positive. Their customers had no idea what was going on. Once it is explained, they are often co-operative. In the abstract, however, the information about risks is not getting through. A security officer at a smaller ISP explained it this way: "At any given point in time, we have 600-800 customers who have a malware, abuse or security problem with their machine. You do not see those numbers in the paper, because a journalist does not think this is a problem; 600 out of 400 000 customers. This is also why end users do not think it is a problem, because the chances of being hit seem so low."

The cost of increasing security provides a further disincentive. The willingness to pay for security services seems low. As quoted earlier, one interviewee summarised their experience as an ISP with offering security software as follows: "If people have the option to pay for it or not to pay for it, they do not." But even after the licence was included in the subscription rate, there was still a large group of people not installing the software package. A similar phenomenon was related to us by the head of Internet security at a large ISP: they too offered an AV solution as part of the subscription. Even the people who did install it often did not keep it up to date. He blamed it on poorly designed software. That sentiment was shared by a representative of a consumer organisation: "We see that the products consumers get for establishing some degree of security for their pc do not work properly and they are too complicated to manage. Consumers cannot manage their own security given the tools they are provided with." When asked whether in their view consumers would be willing to pay for better security, the interviewee responded: "In general terms, they do and they do not. They just expect it to be the default setting. Most products are secure. When you buy a car, it's got seat belts, air bags, brakes. Those things are included in the product. Consumers feel that charging extra for that is a bit ridiculous." In line with these views, a survey by the consumer organisation found that the majority of their members felt that Internet security was a shared responsibility: the consumers themselves are responsible for their online behaviour, but the technical aspects of security are the responsibility of others, most notably their PC retailers, ISPs, software vendors and the government (Consumentenbond 2006).

It is difficult to disentangle incentives from incomplete information, but their combined effect is to undermine the willingness as well as the ability to act. Often this situation is described with a sense of inevitability, as if the home user is a static entity with no learning curve. Surveys suggest that image is incorrect. Home users are adapting their behaviour, but it is unclear how these changes add up, how to connect the disparate if not contradictory pieces of information from the plethora of surveys out there. Even if we ignore the discrepancies between the numbers, it is hard to characterise the current situation. Surveys tell us a large number of people are worried about identity theft, privacy, security, online predators, fraud and other problems. In fact, a significant portion of people are turning away from the Internet altogether (GetSafeOnline 2006). At the same time, adoption of security measures such as firewalls and AV software is increasing, slowly but surely (Fox 2007).

The key question regarding the incentive structure is: how, if at all, are home users confronted with the costs generated by their security tradeoffs? Of course, technically, they are confronted with it all the time. The bulk of the spam messages that everyone receives is sent through botnets, to name but one consequence. But the causality between individual behaviour and such aggregate effects is too abstract and complicated to have a feedback effect. Feedback typically stems from actual security problems that people experience – the victims of fraud, identity theft or, less dramatic, degraded functionality of their machines. According to a 2007 survey of Consumer Reports, 1 in 5 people experience a major virus problem, 1 in 11 experience a major spyware problem and 1 in 81 actually lost money from an account (Consumers Union 2007). Assuming these numbers are correct, that would mean somewhere between 20-30% of all home users have directly experienced the consequences of their security decisions. Potentially, this could be a powerful feedback loop, but the unanswered question is how people understand these incidents, if they relate them back to their own decisions and whether they have adequate tools and capabilities to act on their understanding, assuming such tools exist for end users. The existing security software suites are increasingly ineffective in detecting malware.

The most direct mechanism which is currently internalising some of the externalities generated by end users is the practice of ISPs isolating infected users until they resolve the security problem. It would appear that this solution works for relatively modest numbers of infected machines but, as computer experts say, it does not scale to the actual number of infections. It is not just ISPs that bear the externalities generated by home users. Most online forms of business are confronted with botnets and related security threats and they have to provision their services accordingly – whether this is an e-commerce company buying DDoS

mitigations services from its ISP or an online bank that has to design its services under the – all too valid – assumption that the customer's machine is compromised. Few of these market parties are in a position to mitigate these risks by influence the security tradeoffs of end users. Thus, defending against these security threats is perceived as the cost of doing business.

Large end users

The situation for large organisations – public and private – is rather different. On average, they have dedicated IT staff available and are in a much better position to understand the security risks they face, take precautionary measures as well as build incident response capabilities. Notwithstanding these advantages, research often reports that both public and private organisations underestimate the risks they face or under-invest with regard to security. Some of our interviewees reported compromised machines in their networks, which they perceived as more or less inevitable. They indicated that their networks were by necessity rather open to accommodate contractors or the flexible use of services throughout the organisation. One interviewee said his network was like a fortress which kept intruders out, but once someone had gained a foothold inside, there were many opportunities for malicious activity.

While interviewees reported instances of malware on their network, they claimed this malware to be generic and not targeting their organisation specifically. It is unclear how valid this claim was. The way they found out about these compromised machines – *e.g.* through notification by security service providers which were not under contract with them or during the activities of support desk staff repairing malfunctioning machines – suggests that their risk perception of malware is not based on any formal type of analysis of their own services and networks.

There are many known cases of companies who suffered embarrassing security breaches – and there are undoubtedly many more unknown ones. That being said, it is rather difficult to determine the appropriate level of investment in light of these threats. While more formal analytic instruments have been developed in recent years to support these decisions, their application requires the input of values and probabilities that are very hard to estimate with any degree of reliability. According to the 2007 CSI Computer Crime and Security Survey, less than half of all organisations use instruments such as ROSI, IRR and NPV (CSI 2007). Insurance providers have very little actuarial data to base policies on.

While the security practices of large end users undoubtedly leave much room for improvement, it is also important to realise that many claims that businesses underestimate risks and under-invest in security stem from research that is sponsored or carried out by security providers, with an incentive to overestimate rather than underestimate the problem. Contrast these claims to the findings from the CSI Survey, which published decreasing loss estimates from respondents for five years in a row – a trend only reversed last year (CSI 2007). The peak loss was experienced in 2001 with more than USD 3.1 million per reporting organisation. Since then, most likely due to increased awareness and more systematic investment in computer security, the damages have declined to a low of USD 168 000 per reporting organisation in 2006. In 2007, the downward trend reversed as damages per reporting organisation doubled to USD 345 000. It is difficult to assess whether this represents a one-time deviation or a sustained reversal of the downward trend. Most likely it reflects the technology race between the provision of cybersecurity and ever-more sophisticated and virulent criminal attack techniques. It is also important to note that direct losses are no measure of the complete financial impact felt by society.

Organisations face all kinds of tradeoffs regarding their information security decisions, including malware. Take the issue of patching. We heard estimates that patching mission-critical software systems can cost millions. For that reason, some companies did not patch immediately after release of a vulnerability patch, but waited for months and then applied several patches simultaneously. There were even examples of organisations who consciously never patched, estimating the risk of disruption through

patching to be higher than that of security breaches. In the financial sector, security measures often face a tradeoff against availability of the systems and their performance. In a world where the ability to process information in milliseconds affects the bottom line, measures that improve security but slow down transactions are not an obvious choice. A similar tradeoff exists between security and availability – that is, the uninterrupted uptime of systems. All of these tradeoffs involve difficult assessments of costs and benefits, often in the face of uncertainty and missing information.

Even if it is true that large organisations might not fully understand the costs and benefits of information security, the more relevant issue is whether this situation causes externalities. In the absence of externalities, it is within their purview to pursue whatever security strategy they deem appropriate and bear the consequences of those decisions. In most generic terms, the answer is Yes, there are serious externalities. Hospital records that are compromised, financial records of millions of citizens that are 'lost,' a job website that has been compromised, allowing the personal information of over a million users to be stolen (Wilson 2007). The list goes on and on. If we expand the set of security breaches to also include attacks that did not directly involve malware, the enormous potential for externalities becomes clear. As malware develops and proliferates, it seems reasonable to assume that over time it will be implicated in a wider variety of security breaches than those we have already observed.

What are the incentives for these organisations to prevent these externalities? There is brand damage. Organisations that have been breached have a strong incentive not to disclose this information. However, many US states have adopted legislation that requires organisations to publicly disclose security breaches. The legislation includes no penalties, but still provides strong incentives because of the prospects of public embarrassment and loss of share value. Campbell *et al.* (2003) reported that, on average, breaches of confidentiality had a significant negative impact, causing an average decline of market value of about 5 %. A study by Cavusoglu *et al.* (2004) also reported that announcing an Internet security breach is negatively associated with the market value of the announcing firm. The breached firms in the sample lost, on average, 2.1 % of their market value within 2 days of the announcement — an average loss in market capitalization of USD 1.65 billion per breach. While these effects are significant, some experts argue that these are temporary and that, over time, the notifications will have less and less impact as the number of notifications increases and they lose their news value.

Data breach notification legislation enables other parties to hold the responsible organisation liable for any damages they have suffered. This may be done by individuals affected, but perhaps more realistically by other companies which have more resources to pursue such a course of action. In the case of the security breach at Choicepoint, this led to $10 million in civil penalties for security breaches and USD 5 million in redress to customers (FTC 2006). More recently we have seen what will undoubtedly be a landmark case, the security breach at the US retailer T.J. Maxx in December 2006. Many parties are suing the retailer for damages following this breach. Among them are the banks who reimbursed their own customers for fraudulent transactions using credit card information that was stolen at T.J. Maxx. Recently, the retailer has reported that the breach has already cost them USD 135 million – and the case is far from over. A security company estimated that in the end, it would cost the company around USD 4.5 billion (Gaudin 2007).

US security breach notification laws provide incentives that internalise some of the externalities of security decisions of large organisations. Other US legislation also has implications for liability, most notably Sarbanes-Oxley, the Health Insurance Portability and Accountability Act and the Gramm Leach Bliley Act. While there is disagreement over the effectiveness of these laws, issues of liability and compliance have shown to be drivers for increased security efforts (*e.g.* Ernst & Young 2007; Lords 2007, p. 152). Other countries have different regulatory regimes in place. However, there are parallels. Data protection laws could potentially have similar effects. So far, however, these effects, if they are indeed occurring, are certainly less visible. Predictably, the debate is shifting towards the issue of whether to

connect sanctions to these liabilities. The UK Information Commissioner recently called for criminal sanctions "for those who knowingly and recklessly flout data protection principles" (Shifrin 2007).

Sum

End users have been the focus of considerable debate over Internet security. As has been reported before, many externalities emanate from end users' security decisions – or non-decisions. Interestingly, both for home users and large users, there are incentives which are potentially very strong – that is, the risk of significant damage to themselves resulting directly from their decisions. The problem is, however, that their risk perceptions are often not consistent with the technological realities in which they operate. To the degree that end users do appreciate the risks they face, there are significant problems when acting on that information. For home users, security tools are often too complex and partially effective at best. For large public and private organisations, the situation is remarkably similar. While they often have more expertise available, the security challenges are also substantially more complex in light of the complicated array of systems, services and the organisational arrangements around them. As a result, end users generate externalities, the costs of which are sometimes passed back to them, but which in many cases are internalised by other market players as the cost of doing business in the value net of the information industry, or by society at large.

IV. INCENTIVES AND EXTERNALITIES RELATED TO MALWARE

The economic effects of malware are the joint outcome of the intensity of attacks by illegal and criminal actors and the responses of the players in the affected industries. As the sheer number of new forms of malware and their sophistication increases, the total cost of maintaining a given level of security will, in the absence of significant advances in security technology, increase as well. The overall costs of security as well as the level of security in the value net can be influenced by measures affecting illegal and cyber-criminal activity or via measures reducing the vulnerability of the value net.

The preceding chapter reported on the efforts and incentives of a variety of market players. It indicated a number of market-based incentive mechanisms that contribute to enhanced security but also other instances in which decentralised actions may lead to sub-optimal outcomes. A pressing question is whether the response to malware of actors in information and communication markets is adequate or whether improvements are possible. Pointing to a variety of reports that show increases in malicious attack trends, one might conclude that markets are not responding adequately. Our analysis revealed a more nuanced picture.

Externalities related to malware

According to standard economic theory, an allocation of resources would be efficient if each decision maker properly takes externalities into account. An important claim of welfare economics is that this and other preconditions are met, the overall allocation of resources will be efficient. This may be achieved either because of the intrinsic incentive structure or because of formal or non-formal measures that align the intrinsic incentive structure with the conditions for reaching an overall social optimum. These efficiency conditions are violated repeatedly in the case of malware although the severity varies considerably.

However, several caveats are in place. These optimal conditions are established for fairly abstract and idealised model worlds. Real-world markets rarely meet the preconditions that are assumed to hold. For example, decision makers rarely have complete information, they operate under conditions of bounded rationality, and behave opportunistically. For these reasons, individual decisions rarely are as ideal as described by abstract models. Rather, they are a process of "muddling through" second and third-best solutions, especially in an environment of rapid technological change. Whether a decision was a good or a bad one often is only revealed *ex post* in a dynamic evolutionary process.

Assessing the direct and indirect economic cost of malware under real world conditions is hence an important aspect of designing countermeasures. As the provision of security entails cost, tolerating a certain level of insecurity is economically rational. The level of security realised therefore depends on the costs and benefits of security to individual actors and potential collective measures to enhance security. One key question is whether market players are taking the full range of costs into account when making security decisions. If costs are externalised to other market players or society at large, how serious are they in relation to the internalised costs? While keeping in mind the scope and limitations of our study, we can offer a number of tentative conclusions with regard to these questions. Across the value net of the different market players, three relevant situations emerge:

i) No externalities

This concerns instances in which a decision-making unit, be it an individual user or an organisation, correctly assesses security risks, bears all the costs of protecting against security threats (including those associated with these risks) and adopts appropriate countermeasures. Private and social costs and benefits of security decisions are aligned. There may still be significant damage caused by malware, but this damage is borne by the market player itself. This situation would be economically efficient but, due to the high degree of interdependency in the Internet, it is rare.

That does not mean these situations are non-existent. In principle, end users – be they large organisations or skilled home users – who use adequate security policies and successfully prevent their machines from being compromised generate no externalities for the rest of the value net – though some experts might argue that under certain conditions such behaviour creates positive externalities that are not taken into account and thus lead to an sub-optimal level of private investment (Kunreuther and Heal 2003). Several interviewees claimed that in recent years, they have not had any malware infection within their organisation's network. We were not in a position to check the validity of these claims, but it is not unreasonable to assume that there are cases where malware is successfully fought off or where the effects of malware infections are by and large limited to the owner of the infected system.

ii) Externalities that are borne by agents in the value net that can manage them

This concerns instances in which an individual unit assesses the security risks based on the available information but, due to the existence of (positive or negative) externalities, the resulting decision deviates from the social optimum. Such deviations may be based on lack of incentives to take costs imposed on others into account, but it can also result from a lack of skills to cope with security risks, or financial constraints faced by an individual or organisation. As long as somebody in the value net internalises these costs and this agent is in a position to influence these costs – *i.e.* it can influence the security tradeoffs of the agents generating the externality – then the security level achieved by the whole value net may will deviate less from a social optimum than without such internalisation. This scenario depicts a relatively frequent case and numerous examples were found that confirm externalities were being internalised by other market players.

For example, ISPs have started to manage the security problems generated by their customers – *e.g.* by quarantining the infected machines of customers. As such, they absorb some of the costs generated by the sub-optimally low investment in security by their own customers. ISPs internalise these costs, because not doing would lead to even higher costs being imposed on them, as they may experience blacklisting, rising customer support and abuse management costs and possible reputation effects. The key point here is that ISPs are internalising these costs, but that they are also in a position to influence the behaviour of the agents generating the externality – *i.e.* their own customers. For example, if they increasingly suffer blacklisting because of the spam from infected end user machines going out through their network, one of the options they have is to block port 25. That would significantly reduce the degree of blacklisting and the costs associated with it. Of course, such a measure also has costs and implies a tradeoff with other objectives, such as the kind of services the ISP can offer its customers. They may opt against blocking port 25 for a variety of reasons. That does not take away, however, that the externality is not a given, but that they can actually influence its magnitude. This is different from, say, an e-commerce company who has to buy DDoS mitigation services from its ISP because of botnet attacks. That company can not do anything about botnets and thus the costs to defend itself against them is simply considered a cost of doing business.

ISPs only internalise a part – some experts would say a minor part – of the externalities caused by their customers. For example, while ISPs are increasingly responsive to incoming notifications of abuse on their network, these notifications typically concern only a small fraction of the total number of infected

customer machines. The externalities generated by the remaining machines still affect the wider value net and society at large – see also *iii)*.

Another instance of this type of externality was found in the case of financial services. The incentives of financial service providers are such that in many cases they compensate customers for the damage they suffered from online fraud. In that sense, they internalize the externalities of sub-optimal security investments of their customers as well as the software vendors whose software is exploited to execute the attacks. Many financial service providers claim they compensate all malware-related losses. If that claim is accurate, then the security level achieved by the whole value net may not be too far from the optimum. The financial institutions bear the externalities, but they are also in a position to manage the risk through the security measures around online financial services.

However, there are three important considerations to take into account. First, it is unclear what the reality is of customer compensation under the current liability regime. Some researchers suggest that many claims are in fact refused and that not all of the victim's damage is compensated, only the direct loss (Schneier 2005; Anderson 2007). Second, there is debate within the industry to change the banking codes so as to assign more liability to the customer. New Zealand has already adopted a revised code to this effect. That would change the incentives which might push the level and focus of security investments of the financial institutions away from the social optimum (Anderson 2007, pp. 13-16). Third, even if customer damage is compensated, one could argue that there are still externalities in the sense that important social efficiencies could be gained if people had higher trust in these services and would adopt them more quickly. These benefits would outweigh the additional security investments that would be needed. While the magnitude of these externalities is unknown, the financial service providers are the ones who stand to gain most from maintaining high trust in the e-channel. In other words, this is a problem of incomplete information, rather than of misaligned incentives.

iii) Externalities that are borne by agents who cannot manage them or by society at large

An individual unit may correctly assess the security risks given its perceived incentives but, due to the existence of externalities, this decision deviates from the social optimum. Alternatively, an individual unit may not fully understand the externalities it generates for other actors. Unlike in scenario *ii)*, no other agents in the information and communication value net absorb the cost or, if they do, they are not in a position to influence these costs – *i.e.* influence the security tradeoffs of the agents generating the externality. Hence, costs are generated for the whole sector and society at large. These are the costs of illegal activity or crime associated with malware, the costs of restitution of crime victims, the cost of law enforcement associated with these activities, and so forth. Furthermore, they may take on the more indirect form of slower growth of e-commerce and other activities. Slower growth may entail a significant opportunity cost for society at large if the delayed activities would have contributed to economic efficiency gains and accelerated growth. A comprehensive assessment of these additional costs will demand a concerted effort but will be necessary to determine the optimal level of action to fight malware.

The most poignant cases in this category are the externalities caused by lax security practices of end users. Some of these externalities are internalised by other market players, but many are borne by the sector as a whole and society at large. These externalities are typically explained by the absence of incentives for end users to secure their machines. It would be more precise, however, to argue that the end users do not *perceive* any incentives to secure their machines. While malware writers have purposefully chosen to minimise their impact on the infected host and to direct their attacks at other targets, there is also a plethora of malware which does in fact attack the infected host – most notably to scour any personal information that can be used for financial gain. In that sense, end users do have a strong incentive to secure their machines. Unsecured machines cannot differentiate between malware that does or does not affect the owner of the machine. If the machine is not sufficiently secured, then one has to assume that all forms of

malware can be present. The fact that this incentive is not perceived by the end user is an issue of incomplete information rather than a lack of incentives.

Distributional and efficiency effects

To sum up: Yes, there are significant externalities, but not all of them create sub-optimal solutions. We need to distinguish between distributional and efficiency effects of externalities. When the agent that internalises the externality is in a position to mitigate the risks that generate the externality, then the resulting level of security might not be that far from the social optimum. In this cases, the externality has a distributional effect: it shifts costs (and benefits) from one agent to another. An efficiency effect prevails if a negative or a positive interdependency affects the optimality of decentralized market allocations. Typically, decentralized decision-making will create a deviation of the private from the social optimum in that activities that are afflicted with negative externalities are expanded beyond the optimum level and activities that are afflicted with positive externalities fall short of the optimum level. Like in many other instances, in the case of malware negative externalities are often the flipside of positive externalities: lack of investment by a user in security causes a negative, while investment by a user in security causes a positive externality. The associated effect could be internalised with measures punishing the negative externality effects or measures rewarding the positive aspects.

Distributional aspects need to be considered separately from these efficiency aspects. As discussed, the specific incentive structure of market players often also internalises some of the externalities generated at other stages in the value net. In as far as costs are shifted to another stage of the value net and internalised at that stage, distributional effects prevail. In other words, a mere shifting of the costs (and benefits) between actors takes place. In contrast, overall efficiency gains would materialise if the cost of achieving a given level of information security can be reduced for all the participants in the sector. This differentiation is also important in the evaluation of alternative strategies for coping with problems of malware. Some measures, such as a modification of liability rules, may predominantly shift the burden of combating malware from one set of actors to another. In these cases it will be critical that the resulting attribution of costs and benefits is better aligned with the true cost structure of the value net. Only in this case will efficiency be improved.

Due to the high degree of interrelatedness, nearly all the observable externalities are afflicted with both types of effects. In general terms, however, we would expect that category *ii)* externalities have mainly distributional effects, while category *iii)* will have distributional as well as efficiency effects. From a societal perspective, the latter is obviously a more damaging form of market failure. In the case of *ii)*, efficiency effects are not a given – *i.e.* these cases need not imply a suboptimal level of security for the value net as a whole. Banks, for example, internalise the externalities generated by end users and others. This does not need to have efficiency effects, because the banks can mitigate the risks of end users and thus can tradeoff the damage against the costs of mitigation. In fact, it may have a positive effect on efficiency, if the banks can manage better the risks than the users themselves.

It is important to keep in mind that many malware-related externalities and costs have their origin in illegal and criminal behaviour: not legitimate market players imposing costs on others. In that sense, the oft-cited analogy to externalities in environmental pollution does not hold. In the example of pollution, there is a market player that benefits from the production process causing that pollution. In that case, the guiding principle of standard economic theory is to internalise the costs of pollution so that the agent adjusts the level of production to be more in line with the social optimum. In the case of malware, the agent who profits from the malware is outside the security market. Malware increases the costs of security for all and causes additional direct and indirect costs for damages or foregone activities. As such, it stands to reason that parts of these externalities should be internalised by measures of the sector as a whole or

society at large and not by individual stakeholders. This is currently happening, for example, in the area of law enforcement, but it is not clear whether it is at an optimal level.

With regard to the interrelationships within the information and communications-related activities, it seems that the incentives of many of the commercial stakeholders are reasonably aligned with minimizing the effects of externalities on the sector as a whole. The incentives typically have the correct directionality, but in a variety of cases they are too weak to prevent significant externalities. It is important to note, however, that all market players we studied experience at least some consequences of their security tradeoffs on others. In other words, there was a feedback loop that brought some of the costs imposed on others back to the agent that caused them – even if in some cases, the force of the feedback loop has so far been too weak or too localised to move their behaviour towards more efficient social outcomes.

The costs of malware

Although the malware-related costs of security measures are considered proprietary, estimates provided by players range from 6-10% of the investment in ICT. No clear estimates of the effects of malware on operating expenses were available, although we did find that most organisations did experience such effects. There was evidence throughout the empirical research of concern that such effects are important, although no specific indication as to their magnitude is available. The concern with this broader societal externality seems to motivate several players, especially in industries sensitive to reputation issues, to increase investment in security and to add a "safety margin" when deciding on levels of security.[11]

The total costs of malware include damages experienced by individuals and organisations, private and social costs of preventative and enforcement measures, as well as indirect costs. Private costs of prevention include expenses for increased information security. Social costs of preventative and enforcement measures may include the drafting and enforcement of legal provisions or the cost of developing and enforcing forms of self-regulation. Indirect costs may emanate from missed productivity increases due to lack of trust in electronic transactions.

Assessments of the damages of security incidents reveal a wide variability. We discussed some of these figures in the previous chapter. The respondents to the CSI Computer Crime and Security Survey, widely seen as the best available source of damage estimates, actually reported decreasing security-related losses for five years in a row. Only last year was this trend reversed with an upswing, though still far below the peak losses in 2001 (CSI 2007). Contrast these findings with reports from the financial industry, where online fraud, which includes malware-related losses, has been rising rapidly (Krebs 2008), though the UK financial sector also reported a fall in phishing fraud (APACS 2008). To interpret these numbers appropriately, one would need to relate them to other variables, most notably the overall volume of online transactions. For example, fraud data on credit card use, online and offline, show that losses have increased in absolute terms, but not in relative terms – VISA Europe, for example, recently reported an "all-time low" fraud rate (fraud to sales by cards issued) of 0.051 per cent volume (House of Lords 2007b, p. 36). The available APACS data on online fraud losses is not sufficiently disaggregated to allow separating online transactions from transactions through other channels, though they did claim that the volume of transactions over Internet, phone and mail grew substantially faster than the amount of fraud using these channels (APACS 2008). It is unclear how to evaluate this claim and these trends in light of the available information.

[11] For a literature review of the available estimates of the costs of malware and network security in general, see: Bauer, J. M., M. J. G. Van Eeten and T. Chattopadhyay (Forthcoming). *Financial Aspects of Network Security: Malware and Spam*. ITU (International Telecommunication Union). Available online at: www.itu.int/ITU-D/cyb/.

Direct and indirect costs related to malware are often gauged by mapping the attack trends. While it is true that many attack trends are increasing there is no simple relation between these trends and the overall costs of malware. Detecting a higher number of Trojan variants does not necessarily mean that there is more damage. It could also be a response to improved security defenses or reduced benefits per attack. Similarly, signaling that large-scale botnets are shrinking in size does not necessarily mean that the countermeasures are effective. It might be that attackers have found smaller and more focused botnets to be more profitable. Because malicious attack trends and countermeasures are highly dynamic, it is difficult to draw reliable conclusions on the costs of malware from the attack trends themselves.

The information collected in this research project from actors across the information and communication value net allows the conclusion that the direct private and public costs of prevention are substantial. With few exceptions, many actors have had to increase their security-related investments as a response to the higher benefits of security associated with the types of transactions conducted via the Internet and the increasing number of attacks. However, each actor typically only acts based on the perceived incentives. In literally all cases there were important costs and benefits that accrued at other stages of the value net and were hence outside the decision-making process. Our research showed that due to feedback effects inherent in market co-ordination, the magnitude of these externalities is probably smaller than hitherto assumed. On the other hand many of these externalities remain uncorrected leaving the system overall in a sub-optimal state. The collective costs of fighting malware, ranging from the costs of maintaining public-private organisations such as CERTs or CSIRTs, to the cost of public education campaigns and law enforcement, add to these private costs. Finally, all actors pointed to the potentially high indirect costs of malware in the form of slower migration to efficiency-enhancing forms of electronic transactions. Taken together, the direct and indirect costs of malware could be a double-digit percentage figure of the revenues of players in the information and communication value net.

Although the research in this report was not designed to develop specific policy recommendations, some general concluding remarks are nonetheless offered.[12] We noted that we found many feedback loops which mitigate the externalities arising from security-reducing behaviour. All market players we studied experience such feedback, which potentially brings their tradeoffs closer in alignment with the social optimum. We also noted, however, that in many cases these feedback loops are too weak or localised to effectively change the security tradeoffs from which the externalities emerge. In terms of policy development, a key strategy would be to strengthen the existing feedback loops and create new ones where possible. That would also keep public policy out of the realm of having to decide how secure is secure enough when it comes to defending against malware.

Given the complexity of the interrelationships, there are no panaceas that could address all the issues with one sweep. From our analysis we conclude that measures increasing the costs of illegal and criminal activities will, other things being equal, help reduce the overall costs of security but, as actors may reduce their investments in security, they may not necessarily increase the overall level of security. If this were the case, such measures, while reducing the direct costs of security, may not reduce the costs of damages associated with security breaches. On the other hand, positive and negative measures increasing the level of security may increase the cost of security but will also have clear negative effects on associated damages, again all other things being equal. In a highly interrelated system, it is often difficult to assess the overall impact of a policy measure due to feedback and unanticipated effects. It is therefore necessary to search for measures that are robust and have desired overall effects in multiple scenarios. In many cases this may require a clarification of the rights and obligations of individuals or classes of stakeholders.

[12] For those readers interested in policy recommendations, we point to a recent study: 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.

REFERENCES

A-i3 (2006). Zur Haftung von Phishing-Opfern. *Arbeitsgruppe Identitätsschutz im Internet e.V.* Available online at https://www.a-i3.org/content/view/975/230/.

Anderson, R. (2001). *Why Information Security is Hard: An Economic Perspective*. Proceedings of the 17th Annual Computer Security Applications Conference, New Orleans, Louisiana IEEE Computer Society. Available online at http://www.acsac.org/2001/papers/110.pdf.

Anderson, R. (2002). *Maybe we spend too much? Unsettling Parallels Between Security and the Environment*. First Annual Workshop on Economics and Information Security online at http://www.cl.cam.ac.uk/~rja14/econws/37.txt.

Anderson, R. (2007). *Closing the Phishing Hole – Fraud, Risk and Nonbanks*. online at http://www.cl.cam.ac.uk/~rja14/Papers/nonbanks.pdf.

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

Anderson, R. and T. Moore (2006). The Economics of Information Security. *Science* 314: 610-613.

Anderson, R. and T. Moore (2007). *Information Security Economics – and Beyond*. Computer Laboratory, University of Cambridge. Available online at http://www.cl.cam.ac.uk/~rja14/Papers/econ_crypto.pdf.

APACS (2008). *Press release: Fraud abroad pushes up losses on UK cards following two-year fall*. Available online at http://www.apacs.org.uk/2007Fraudfiguresrelease.html.

Arbor Networks (2007). *Worldwide Infrastructure Security Report, Volume III*. Online at http://www.arbornetworks.com/report.

August, T. and T. I. Tunca (2006). Network Software Security and User Incentives. *Management Science* 52(11): 1703–1720.

Bangeman, E. (2006). *Court likely to order ICANN to suspend Spamhaus' domain*. Ars Technica. Available online at http://arstechnica.com/news.ars/post/20061009-7938.html.

Banktip (2006). *Phishing: Kunden haften für Trojaner*. Banktip.de. Available online at http://www.banktip.de/News/20648/Phishing-Kunden-haften-fuer-Trojaner.html.

Barnum, S. and M. Gegick (2005). *Economy of Mechanism*. Build Security In. Available online at https://buildsecurityin.us-cert.gov/daisy/bsi/articles/knowledge/principles/348.html?branch=1&language=1.

Bauer, J. M., M. J. G. Van Eeten and T. Chattopadhyay (Forthcoming). *Financial Aspects of Network Security: Malware and Spam*. ITU (International Telecommunication Union). Available online at http://www.itu.int/ITU-D/cyb/.

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.

Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *The Journal of Political Economy* 76(2): 169-217.

Becsi, Z. (1999). Economics and Crime in the States. *Economic Review - Federal Reserve Bank of Atlanta* 84(1): 38-56. Available online at http://ezproxy.msu.edu:2047/login?url=http://proquest.umi.com/pqdweb?did=40779835&Fmt=7&clientId=3552&RQT=309&VName=PQD

Berner, R. and A. Carter (2005), "The truth about credit-card fraud", *Business Week Online*. Available online at http://www.businessweek.com/technology/content/jun2005/tc20050621_3238_tc024.htm.

Bernstein, D. J. (2007). *Some thoughts on security after ten years of qmail 1.0*. 1st Computer Security Architecture Workshop in conjunction with 14th ACM Conference on Computers and Communication Security, Fairfax, Virginia. Available online at http://cr.yp.to/qmail/qmailsec-20071101.pdf.

Böhme, R. (2005). *Cyber-Insurance Revisited*. Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/15.pdf.

Camp, L. J. (2006). *Mental Models of Privacy and Security*. Available online at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=922735.

Camp, L. J. and C. Wolfram (2004). *Pricing Security: Vulnerability as Externalities*. Available online at http://ssrn.com/abstract=894966.

Campbell, K., L. A. Gordon, M. P. Loeb and L. Zhou (2003). The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market. *Journal of Computer Security* 11(3): 431-448. Available online at http://brief.weburb.dk/archive/00000130/01/2003-costs-security-on-stockvalue-9972866.pdf.

Cavusoglu, H., H. Cavusoglu and S. Raghunathan (2005). *Emerging issues in responsible vulnerability disclosure* Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/cavusoglu.pdf.

Cavusoglu, H., B. Mishra and S. Raghunathan (2004). The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers. *International Journal of Electronic Commerce* 9(1): 69. Available online at http://www.gvsu.edu/business/ijec/v9n1/p069.html.

Chen, P.-Y., G. Kataria and R. Krishnan (2005). *Software Diversity for Information Security*. Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/47.pdf.

Choi, J. P., C. Fershtman and N. Gandal (2005). *Internet Security, Vulnerability Disclosure, and Software Provision*. Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/9.pdf.

Clayton, R. (2007). *Phishing and the gaining of "clue"*. Available online at http://www.lightbluetouchpaper.org/2007/08/16/phishing-and-the-gaining-of-clue/.

Consumentenbond (2006). PC beveiliging & veilig Internet: Een enquête onder computergebruikers. *Consumentengids* 2006(11).

Consumers Union (2007). 'State of the 'Net' Survey '07. *Consumer Reports* 2007(9): 28-34.

Counterpane & MessageLabs (2006). *2005 Attack Trends & Analysis*. Available online at http://www.counterpane.com/dl/attack-trends-2005-messagelabs.pdf.

CSI (2007). *CSI Survey 2007: The 12th Annual Computer Crime and Security Survey*. Computer Security Institute. Available online at http://www.gocsi.com/forms/csi_survey.jhtml.

Dot-TK (2007). *Dot Tk Free Domain Names – A New Approach To Make A Whole Top Level Country Domain Free Of Illicit Content*. Available online at http://www.dot.tk/en/press_jul16-07.pdf.

Dynes, S., E. Andrijicic and M. E. Johnson (2006). *Costs to the U.E. Economy of Information Infrastructure Failure from Field Studies and Economic Data*. Fifth Workshop on the Economics of Information Security 2006, http://weis2006.econinfosec.org/docs/4.pdf. Available online at http://weis2006.econinfosec.org/docs/4.pdf.

Dynes, S., H. Brechbühl and M. E. Johnson (2005). *Information Security in the Extended Enterprise: Some Initial Results From a Field Study of an Industrial Firm*. Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/51.pdf.

Ehrlich, I. (1996). Crime, Punishment, and the Market for Offenses. *The Journal of Economic Perspectives* 10(1): 43-67. Available online at http://links.jstor.org/sici?sici=0895-3309%28199624%2910%3A1%3C43%3ACPATMF%3E2.0.CO%3B2-U.

ENISA (2006). *Provider Security Measures Part 1: Security and Anti-Spam Measures of Electronic Communication Service Providers - Survey*. European Network and Information Security Agency. Available online at http://www.enisa.europa.eu/doc/pdf/deliverables/enisa_security_spam.pdf.

Ernst & Young (2007). *Global Information Security Survey 2006*. Available online at http://www.ey.nl/download/publicatie/2006_GISS_EYG_AU0022.pdf.

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

Franklin, J., V. Paxson, A. Perrig and S. Savage (2007). *An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants*. CCS'07. Available online at http://www.icir.org/vern/papers/miscreant-wealth.ccs07.pdf.

Friedman, L. S. (2002). *The Microeconomics of Public Policy Analysis*. Princeton, NJ, Princeton University Press.

FTC (2006). *ChoicePoint Settles Data Security Breach Charges; to Pay $10 Million in Civil Penalties, $5 Million for Consumer Redress*. Available online at http://www.ftc.gov/opa/2006/01/choicepoint.shtm, November 25,2007.

Gal-Or, E. and A. Ghose (2003). *The Economic Consequences of Sharing Security Information*. 2nd Annual Workshop on Economics and Information Security. Available online at http://www.cpppe.umd.edu/rhsmith3/papers/Final_session7_galor.ghose.pdf.

Gal-Or, E. and A. Ghose (2005). The Economic Incentives for Sharing Security Information. *Information Systems Research* 16(2): 186-208. Available online at http://www.andrew.cmu.edu/user/aghose/Infosec.pdf.

Gaudin, S. (2007). *T.J. Maxx Security Breach Costs Soar To 10 Times Earlier Estimate*. Information Week. Available online at http://www.informationweek.com/shared/printableArticle.jhtml?articleID=201800259.

GetSafeOnline (2006). *The Get Safe Online Report*. online at http://www.getsafeonline.org/media/GSO_Cyber_Report_2006.pdf.

Gordon, L. A. and M. P. Loeb (2002). The Economics of Information Security Investment. *ACM Transactions on Information and System Security*. Available online at http://portal.acm.org/citation.cfm?id=581274.

Heidrich, J. (2007). *IP-Blacklisting zur Spam-Abwehr kann rechtswidrig sein*. Heise Online. Available online at http://www.heise.de/newsticker/meldung/97568.

Higgins, K. J. (2007a). *Battling Bots, Doing No Harm*. Dark Reading. Available online at http://www.darkreading.com/document.asp?doc_id=118739.

Higgins, K. J. (2007b). *Untying the Bot Knot*. Dark Reading. Available online at http://www.darkreading.com/document.asp?doc_id=114081&WT.svl=news1_6.

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

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

Just, R. E., D. L. Hueth and A. Schmitz (2004). *The Welfare Economics of Public Policy: A Practical Approach to Project and Policy Evaluation*. Cheltenham, UK and Northampton, MA, Edward Elgar.

Krebs, B. (2007). *Study: $3.2 Billion Lost to Phishing in 2007*. Washington Post Security Fix weblog. Available online at http://blog.washingtonpost.com/securityfix/2007/12/study_32_billion_lost_to_phish_1.html.

Krebs, B. (2008). *Banks: Losses From Computer Intrusions Up in 2007*. Washington Post Security Fix weblog. Available online at http://blog.washingtonpost.com/securityfix/2008/02/banks_losses_from_computer_int.html.

Kunreuther, H. and G. Heal (2003). Interdependent security. *Journal of Risk and Uncertainty* 26(2): 231.

Lacohée, H., S. Crane and A. Phippen (2006). *Trustguide: Final Report*. BT Group Chief Technology Office, Research & Venturing / HP Labs / University of Plymouth, Network Research Group. Available online at http://www.trustguide.org.uk/Trustguide%20-%20Final%20Report.pdf.

LaRose, R., N. Rifon, S. Liu and D. Lee (2005). *Understanding Online Safety Behavior: A Multivariate Model*. International Communication Association New York. Available online at http://www.msu.edu/~isafety/papers/ICApanelmult21.htm.

Lemos, R. (2006). *Attackers pass on OS, aim for drivers and apps*. SecurityFocus. Available online at http://www.securityfocus.com/news/11404.

Lords, H. o. (2007). *Science and Technology Committee, 5th Report of Session 2006–07, Personal Internet Security, Volume II: Evidence*. London, Authority of the House of Lords.

Marshall, A. (1920). *Principles of Economics: An Introductory Volume*. London, Macmillan.

Mell, P., K. Kent and J. Nusbaum (2005). *Guide to Malware Incident Prevention and Handling*. National Institute of Standards and Technology. Available online at http://csrc.nist.gov/publications/nistpubs/800-83/SP800-83.pdf.

Microsoft (2005). *The Trustworthy Computing Security Development Lifecycle*. online at http://msdn2.microsoft.com/en-us/library/ms995349.aspx.

Microsoft (2007). *Storm Drain*. Anti-Malware Engineering Team Weblog. Available online at http://blogs.technet.com/antimalware/archive/2007/09/20/storm-drain.aspx.

Nowotny, E. (1987). *Der öffentliche Sektor: Einführung in die Finanzwissenschaft*. Berlin, Springer.

OECD (2002a). *Guidelines for the Security of Information Systems and Networks*. online at http://www.oecd.org/dataoecd/16/22/15582260.pdf.

OECD (2002b). *The OECD 2002 Security Guidelines - Q&A*. online at http://www.oecd.org/dataoecd/27/6/2494779.pdf.

OECD (2005). *The Promotion of a Culture of Security for Information Systems and Networks in OECD Countries*. Available online at http://www.oecd.org/dataoecd/16/27/35884541.pdf.

ORF (2007). *Spamhaus antwortet auf nic.at*. futurezone.ORF.at. Available online at http://futurezone.orf.at/it/stories/201738/.

PayPal (2007). *Key Financial Facts*. Available online at http://www.pppress.co.uk/Content/Detail.asp?ReleaseID=5&NewsAreaID=22&SearchCategoryID=-1.

Pigou, A. C. (1932). *The Economics of Welfare*. London, Macmillan.

Poindexter, J. C., J. B. Earp and D. L. Baumer (2006). An experimental economics approach toward quantifying online privacy choices. *Information Systems Frontiers* 8(5): 363-374.

Rescorla, E. (2004). *Is finding security holes a good idea?* . Workshop on Economics and Information Security 2004. Available online at http://www.rtfm.com/bugrate.pdf.

Rifon, N., E. T. Quilliam and R. LaRose (2005). *Consumer Perceptions of Online Safety*. International Communication Association New York. Available online at http://www.msu.edu/~isafety/papers/ICApanelfg.htm.

Rowe, B. R. and M. P. Gallaher (2006). *Private Sector Cyber Security Investment: An Empirical Analysis*. Fifth Workshop on the Economics of Information Security, 2006, , Cambridge, UK. Available online at http://weis2006.econinfosec.org/docs/18.pdf.

Schechter, S. E. (2004). *Computer Security Strength & Risk:A Quantitative Approach*. The Division of Engineering and Applied Sciences. Harvard University. Available online at http://www.eecs.harvard.edu/~stuart/papers/thesis.pdf.

Schneier, B. (2000). *Secrets and Lies: Digital Security in a Networked World*. New York, John Wiley.

Schneier, B. (2005). "A Real Remedy for Phishers". *Wired News*. Available online at: www.wired.com/news/politics/0,1283,69076,00.html.

Schneier, B. (2007). *Information Security and Externalities*. NSF/OECD Workshop on Social & Economic Factors Shaping The Future Of The Internet, Washington, DC. Available online at http://www.oecd.org/dataoecd/60/8/37985707.pdf.

Shifrin, T. (2007). *Lose an unencrypted laptop and 'face criminal action'*. Computerworld UK. Available online at http://www.computerworlduk.com/management/security/data-control/news/index.cfm?newsid=6241.

Shostack, A. (2005). *Avoiding Liability: An Alternative Route to More Secure Products*. Fourth Workshop on the Economics of Information Security, Harvard University. Available online at http://infosecon.net/workshop/pdf/44.pdf.

Snyder, W. (2007). *Time to Deploy improvement of 25 %*. Mozilla Security Blog. Available online at http://blog.mozilla.com/security/2007/06/18/time-to-deploy-improvement-of-25-percent/.

Sokolov, D. A. (2007). *Spamhaus.org setzt Österreichs Domainverwaltung unter Druck*. Available online at http://www.heise.de/newsticker/meldung/91417.

South, G. (2007). Web issues over banking code. *The New Zealand Herald*. Available online at http://www.nzherald.co.nz/topic/story.cfm?c_id=126&objectid=10458545.

Spamhaus (2007). *Report on the criminal 'Rock Phish' domains registered at Nic.at*. Available online at http://www.spamhaus.org/organization/statement.lasso?ref=7.

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 http://www.bsi.de/literat/studien/recht/Gutachten.pdf.

TechWebNews (2005). "One In Four Identity-Theft Victims Never Fully Recover". *Information Week*. Available online at http://www.informationweek.com/showArticle.jhtml?articleID=166402700.

Utter, D. (2007). *MySpace Asked GoDaddy To Drop SecLists*. SecurityProNews. Available online at: http://www.securitypronews.com/insiderreports/insider/spn-49-20070126MySpaceAskedGoDaddyToDropSecLists.html.

Vijayan, J. (2003). Improved security through IT diversity. *Computerworld* 37(47): 28. Available online at http://www.computerworld.com/printthis/2003/0,4814,87470,00.html.

Watson, A. (2005). *Reputation in Open Source Software*. Available online at: http://opensource.mit.edu/papers/watson.pdf.

Weber, T. (2007). *Criminals 'may overwhelm the web'*. BBC News website. Available online at http://news.bbc.co.uk/2/hi/business/6298641.stm.

Wilson, T. (2007). *Trojan on Monster.com Steals Personal Data*. Dark Reading. Available online at http://www.darkreading.com/document.asp?doc_id=131953.

APPENDIX: LIST OF INTERVIEWEES

Alhadeff, Joseph	Oracle [US]
Barbir, Suzana	Telstra BigPond [AUS]
Barrett, Michael	PayPal [US]
Beale, Jeremy	Confederation of British Industry [UK]
Behlendorf, Brian	Mozilla Foundation [US]
Boudewijns, Arno	St. Elisabeth hospital [NL]
Butler, Ben	Go Daddy [US]
Candel, Hans	St. Elisabeth hospital [NL]
Davidson, Mary Ann	Oracle [US]
Dupon, Koen	Consumentenbond (Consumers Union) [NL]
Edelstein, Eric	France Telecom / Orange [FR]
Florijn, Gert	ABN AMRO [NL]
Gorbutt, John	StreamShield [UK]
Hafkamp, Wim	FI-ISAC / Rabobank [NL]
Halfweeg, Jaap	KPN [NL]
Hania, Simon	XS4All [NL]
Hiskey, Steve	Microsoft [US]
Kelly, John	Comcast [US]
Keogh, Steve	Telstra BigPond [AUS]
Lappas, Paul	ServePath [US]
Leguit, Douwe	GOVCERT [NL]
Lord, Peter	Oracle [US]
McIntyre, Scott	XS4All [NL]
Melein, Johan	SIDN (Foundation for Internet Domain Registration) [NL]
Mitchell, Alan	IBM [US]
Molenaar, Danyel	OPTA [NL]
Morrow, Chris	Verizon Business [US]
O'Donnell, Adam	Cloudmark [US]
Oppenheimer, Jay	Comcast [US]
Pinkney , Graeme	Symantec [UK]
Piscitello, Dave	Fellow to the ICANN SSAC [US]
Provos, Niels	Google [US]
Quaresima, Richard	Federal Trade Commission [US]
Rader, Ross	Tucows [CA]
Ramsauer, Thomas	BSI (Federal Office for Information Security) [DE]
Rand, Dave	TrendMicro [JP]
Reed, Chris	Queen Mary University of London [UK]
Reijers, Roeland	GOVCERT [NL]
Renten, Jerry	KPN [NL]
Salsburg, Daniel	Federal Trade Commission [US]
Samson, Michael	NVB (Dutch Association of Banks) [NL]
Schindler, Werner	BSI (Federal Office for Information Security) [DE]

Schoen, Kevin	ACDNet [US]
Schuurman, Jacques	Surfnet CERT [NL]
Silversin, Louis	Federal Trade Commission [US]
Slim, Arjen	Shell International [NL]
Truman, Nick	BT [UK]
Van Daalen, Frits	ABN AMRO [NL]
Van der Heide, Martijn	KPN-CERT [NL]
Veysset, Franck	France Telecom / Orange [FR]
Walsh, Anthony	Shell International [NL]
Ward, Jeremy	Symantec [UK]
Wesson, Rick	Support Intelligence / Alice's registry [US]
Whitaker, Colin	APACS [UK]
Wiggins, Rich	Michigan State University [US]
Williams, Jeff	Microsoft [US]
Woodcock, Bill	Packet Clearing House [US]