Securing Cyberspace:
Realigning Economic Incentives in the ICT Value Net

Johannes M. Bauer and Michel van Eeten
Michigan State University and TU Delft
WebSci'09, Athens, Greece, March 18-20, 2009

1. Introduction

Malicious software ("malware") has become a serious security risk to all users of the Internet, whether they are large or small organizations or home users. Viruses, worms and numerous other variants of malware have developed from a nuisance to sophisticated tools for criminals. Computers all across the world, some estimate as many as 1 in 10 to 1 in 5, are infected with malware without knowledge of the machine's owner. While hacking continues to be one approach, various other infection strategies are widespread. They range from attachments, downloads from websites, the distribution via social networks and games, and reliance on mobile data carriers. Infected machines may be connected in botnets: flexible remote-controlled networks of computers that operate collectively to provide a platform for criminal and fraudulent purposes. Attacks target individuals, organizations, and nations states Spam (most originating from botnets), variations of socially engineered fraud such as phishing and whaling, identity theft, attacks on websites, corporate, "click fraud", "malvertising", and corporate espionage ar some of the attack vectors directed toward individuals and organizations. The massive DDoS attacks on Estonia in 2007[1] and the spread of the cryptic Conficker worm that, early in 2009, paralyzed parts of the British and French military as well as government and health institutions in other countries, are recent examples of attacks on nations and their civil and military infrastructures.[2]

Estimates as to the total cost to society of information security breaches vary widely. A plethora of numbers is available for individual firms and specific sectors such as the financial industries. From a societal point of view, not only the direct cost (e.g., repair cost, losses due to fraud) but also indirect costs (e.g., costs of preventative measures) and implicit costs (e.g., slower productivity increases due to reduced trust in electronic transactions) would have to be attributed to information security breaches. Bauer, Van Eeten, Chattopadhyay and Wu (2008), reviewing a broad range of other studies, conclude that a conservative estimate of these cost may be in the range of 0.2-0.4% of GDP. A catastrophic security failure in the world information and communication system with potentially much higher impact on the global economy is possible. That this is a low-probability event complicates effective preparation and creates the danger of ignoring it altogether.

The analytical lenses of scholars and practitioners on malware have changed significantly in recent years. Initially, security threats were viewed predominantly as technological problems. In recent years, the perspectives broadened to include the economic incentive structures of players in the information and communication technology (ICT) value net and user behavior. This paper focuses on the second aspect by adopting an economic approach to examine the relevant incentives of stakeholders in the (ICT) value net to provide for security and the consequences of

[1] See N. Anderson, "Massive DDoS attacks target Estonia; Russia accused," *Ars Technica*, May 14, 2007, http://arstechnica.com/security/news/2007/05/massive-ddos-attacks-target-estonia-russia-accused.ars# (last visited February 28, 2009).
[2] See M. E. Soper, "Conficker work shuts down French and UK Air Forces," *Maximum PC*, February 10, 2009, http://www.maximumpc.com/article/news/conficker_worm_shuts_down_french_and_uk_air_forces (last visited February 28, 2009).

these incentives for other players and the sector as a whole. Moreover, it takes into account the interdependence between the underground market for cybercrime and Internet security issues. The approach recognizes that much of the problem originates in criminal behavior but it also acknowledges that the magnitude and impact of the malware threat is influenced by the decisions and behavior of legitimate market players, such as Internet Service Providers (ISPs), software vendors, e-commerce companies, hardware manufacturers, registrars and, last but not least, end users. All of these participants in the ICT value net are confronted with malware, but in very different ways. Their capability and effectiveness in counteracting threats and consequently the costs and benefits of differing responses to malware vary widely. In other words, they operate under differing incentives which will affect their individual choices.

As information security comes at a cost, tolerating some level of insecurity is economically rational from an individual and social point of view. From a societal perspective, a key question is whether the costs and benefits perceived by market players are aligned with the social costs and benefits. In that case, individual decisions also ascertain an overall desirable outcome. However, if individual security decisions do not properly reflect social benefits and costs this conclusion does not hold. The security decisions of a market player regarding malware may be rational for that player, given the costs and benefits it perceives. But if the resulting course of action inadvertently or deliberately imposes costs (or benefits) on other market players or on society at large, these are typically not taken into account by a profit-oriented market player. They constitute, in economic terms, an "externality", an interdependence that is not reflected in the cost and benefit calculus of a decision-maker. As externalities are not compensated in market transactions, their presence leads to sub-optimal outcomes of decentralized decision-making processes, constituting a form of market failure. In the case of information security the presence of externalities may result in internet-based services that are less secure than is socially desirable. As the problem is unlikely overcome by decentralized decisions, it may require collective action in the form of private (voluntary or government-mandated) cooperation and/or public policies to address these shortfalls.

The paper sets out to identify externalities by analyzing the incentives under which a variety of market players operate when dealing with malware. The core of the paper is a detailed discussion of the findings of a qualitative empirical field study. In the course of 2007, a team of researchers from the Delft University of Technology and Michigan State University conducted 41 in-depth interviews with 57 professionals of organizations operating in networked computer environments that are confronted with malware. Interviewees represented a stratified sample of professionals from different industry segments (e.g., hardware, software, service providers, and users) in six countries (Australia, Germany, the Netherlands, United Kingdom, France, and the United States). Moreover, we interviewed experts involved in the governance of information security issues such as Computer Emergency Response Teams (CERTs) and regulatory agencies. Based on this unique and rich data, we identified and analyzed the consequences of the incentives relevant for key players.[3] The next section of the paper discusses the co-evolution of the realms of cybercrime and information security, adopting a market framework. Section three discusses the findings from the empirical study and section four identifies emerging patterns and typical security scenarios. Section five examines policy implications and the final section recaps the main points of the paper.

[3] See Van Eeten and Bauer (2008) for the full report and OECD (2009), in particular Part II.

2. The co-evolution of cybercrime and information security

Information security breaches have become increasingly driven by financial motives. The particular challenges to achieving a desirable level of information security can be better understood by analyzing the interplay of two sets of markets: the market for cybercrime and market for information security. These two markets are tightly linked; developments in one directly affect conditions in the other. It may seem frivolous to think of cybercrime as a "market" but many players in this increasingly differentiated realm do respond to price signals and other economic incentives. (Some players, whether they are inspired by essentially laudable goals (such as "white hat" hackers) or destructive motives (such as cyberterrorists) do not primarily follow financial gain but their decisions may be modelled as an optimization over non-financial goals.) If the logic of these transactions were better known, it might be possible to interrupt and manipulate price signals in ways that quench illegal activity. Very little systematic knowledge is available about the underground cybercrime economy (Schipka 2007; Jakobsson and Zulfikar 2008; Kanich, Kreibnich et al. 2008). More information is available on the market for information security but even this information is often kept proprietary (e.g., CSI 2008). Attack and defence strategies can be analyzed at the level of individual players (e.g., a cybercriminal, a firm investing in information security, a home user deciding on the purchase of a firewall) or at an aggregate (sector) level. In the latter case the interrelationships between players and the effects of their decisions on the working of the value net and the sector as a whole are of primary interest (see Figure 1).

Figure 1: The ICT value net

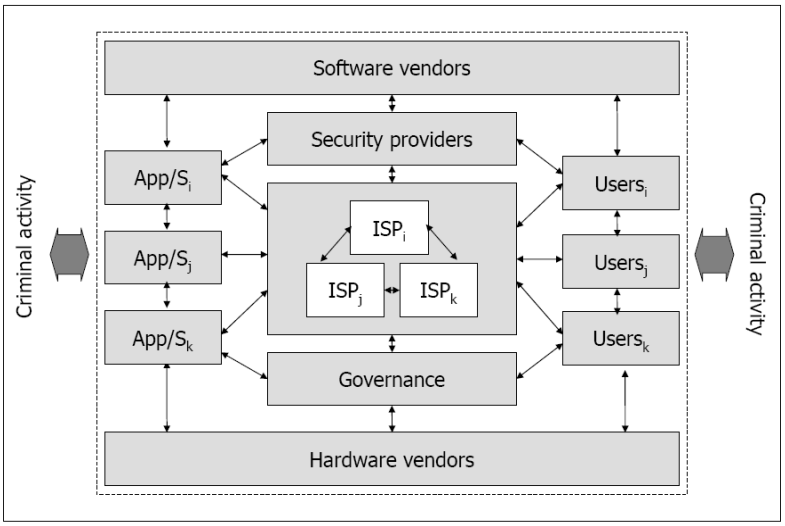

Source: Van Eeten and Bauer (2008)

With the growth of the underground cybercrime economy, different tasks in this network of transactions have become increasingly specialized. One individual or organization is rarely involved in the whole range of activities necessary. The division of labor has led to a differentiation of players such as writers of malicious code, distributors of code, herders ("owners") of botnets, spammers, identity collectors, drop services, drop site developers, and drops (Schipka 2007). Specialization has not only increased the sophistication and virility of malware, it has also increased the productivity of the cybercrime economy and hence reduced the

costs of attack tools. Despite the heterogeneity of these actors, it is reasonable to assume that they act purposefully rational. That is, given their information about relevant opportunities and constraints they will pursue and expand their activities as long as incremental benefits exceed incremental costs. Too little information is typically available to determine the exact shape of cost and benefit relations of the players empirically. However, it is possible to derive some insights from a conceptual analysis. Other things equal ("ceteris paribus"), it is likely that each actor can only expand his or her activity level at higher cost. For example, after highly vulnerable information systems are attacked, it will be exceedingly difficult to penetrate more secure systems. Likewise, it will generally be more time-consuming and hence costly to write code to attack software and devices that have been fortified against attacks. Hence, the short-run incremental cost curve faced by criminals will most likely be upward sloping.

The shape of the incremental benefit relationship is less obvious. It is not possible to reason a priori that for any given type of attack the bigger rewards are easier to reap, implying that the incremental benefit curve is downward sloping. It could well be upward sloping if more effort is "rewarded" with higher spoils. Although this is possible in some cases, it seems unlikely that it can hold across the board, at least in societies that adhere to the rule of law. For the purposes of our further discussion we will therefore focus on the scenario in which the slope of the incremental cost curve is higher than that of the incremental benefits curve. Individual decisions can be aggregated to represent the activities in a specific "market segment", for example, the market for malicious code, botnets, or stolen credit cards. There is no reason to believe that these market segments do not exhibit upward slowing supply and downward sloping demand schedules. Cybercriminal activity could be aggregated even further to generate a representation of the overall market for cybercrime. To this end, it is necessary to define the traded services in a more abstract way. In this aggregate market, supply is shaped by the cost of breaching information systems. Demand is shaped by the benefits of such breaches and the corresponding willingness to pay for them. The intersection of these two schedules will determine the overall level of fraudulent and criminal security breaches (see Figure 2). As long as the incremental benefits in this market exceed the incremental costs, cybercriminal activity will expand (and vice versa).

Figure 2: Markets for cybercrime and cybersecurity

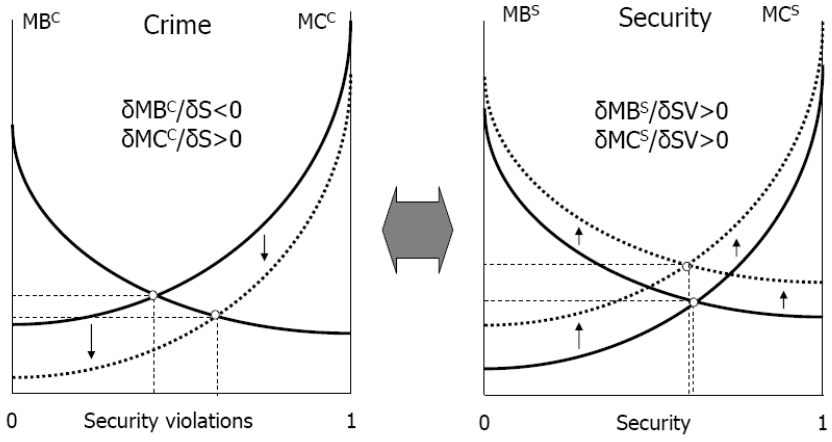

Source: Van Eeten and Bauer (2008)

The shape and position of the schedules is influenced by changes in the technological basis of information systems, the technology of attacks, and the institutional and legal framework governing information security and its violations. Higher penalties for cybercrime, stricter law enforcement that elevates the risk of being caught, and measures that increase information security all shift the supply curve to the upper left. As a consequence, a given level of criminal activity will be associated with higher cost. For a given demand curve for cybercrime (i.e., ceteris paribus) this implies that cybercriminal activity will be reduced. In contrast, a deeper division of labor in the underground economy, improvements in attack technologies, or less diligent law enforcement will shift the supply curve to the lower right. Accordingly, for a given demand curve for cybercrime (i.e., ceteris paribus) cybercriminal activity will increase. The demand schedule for cybercrime is also influenced by other forces. It shifts to the upper right as the net benefits of criminal activity increase, for example, if the number of users of electronic transactions (and hence the potential rewards from cybercrime) grow, connectivity is increasing globally, mobile devices are used more, and ICT use in the private and public sectors becomes more widespread. Ceteris paribus this will lead to increased cybercriminal activity. On the other hand, if the net rewards of cybercrime shrink, for example, because of a lower user base, tighter security, and more stringent law enforcement, the schedule is shifted to the lower left, reducing cybercrime. Presently, the factors contributing to an intensification of cybercrime seem to outweigh the others, putting upward pressure on the overall level of cybercrime.

Like the realm of cybercrime, information security can be analyzed in a market framework. It is best seen as an aggregate made up of individual sub-markets in which vendors supply and heterogeneous types of users demand information security products and services. Such services are offered to different users by a range of specialized vendors, including security service providers (anti-malware software providers, suppliers of network monitoring hardware and software such as Cloudmark), Internet Service Providers, or they may be provided in-house by a specialized department. These players are supported by non-profit groups like Spamhaus, the Messaging Anti-Abuse Working Group (MAAWG), and many blacklisting services that monitor spam and other malicious activity. Security services are ultimately demanded by different types of users including residential users, for-profit and not-for profit organizations of different size, and government agencies. Within its specific context, incentives, and available information, each decentralized player will again make purposefully rational decisions. Once all adjustments are made, it is reasonable to assume that each decision-maker will strive for a situation in which the perceived incremental benefits of additional security measures are approximately equal to the incremental costs of such measures. Given the state of security technology, the incremental cost of improved security will most likely increase. At the same time, with very few if any exceptions, the benefits of additional security measures will decrease. Consequently, the optimal level of security is found where the social cost and benefits of security are equated. In a dynamic perspective, under conditions of risk and uncertainty, although the analysis is somewhat more cumbersome, the same principal decision rule applies: that the optimal level of security is found where the adjusted and discounted benefits outweigh the costs (Gordon and Loeb 2004).

Many of the challenges of reaching an optimal level of information security at the aggregate level are rooted in a potential mismatch between the perceived individual and social benefits and costs of security. In information systems, positive and negative externalities are closely intertwined aspects of security decisions. Additional security investments of end users or of an intermediate service provider such as an ISP also benefit others in the ICT system and are thus associated with positive externalities. Similarly, insufficient security investment exerts a negative externality on others in the value net. The externality problem is complicated by the economic features of information and security markets. Many information products are produced under conditions of high fixed but low incremental costs. Unfettered competition will drive prices quickly to incremental costs. Moreover, many products and services are characterized by positive

network effects. Firms will respond to such conditions with product and price differentiation, attempts to capitalize on first mover advantages, versioning, and other strategies to overcome these challenges (Shapiro and Varian 1999; Varian, Farrell et al. 2005; Illing and Peitz 2006).. Given the current legal framework of information industries, in particular the current liability regime, security is a quasi-public good. Suppliers will typically not be able to mark-up their products in accordance with the social value of their security features but only with the private value, leading to an under-provision of security. If suppliers face a trade-off between speed to market and security, given positive network effects, speed to market may take precedence over security testing (Anderson 2001; Anderson and Moore 2006).

Cybercriminal activity and decisions in information security markets co-evolve, mutually influencing but not fully determining each other's course. An increased level of cybercrime (determined in the market for cybercrime) will increase the effort and cost of maintaining a given level of information security. At the same time, the benefits of increased security may increase. Consequently, the overall cost of security will increase for the participants in the whole ICT value net even if the level of security remains unchanged. In contrast, if effective measures reduce the level of cybercrime, the costs of security will decline as well. However, the overall level of security may not increase as users decide to maintain a given level with lower expenditures. There is an asymmetry in the relation between the markets for cybercrime and cybersecurity. Whereas higher or lower levels of cybercrime will increase or decrease the overall cost of security, the effect of such changes on the level of security remains ambiguous. In contrast, increased security will unequivocally increase the cost of cybercrime and at the same time reduce the benefits of cybercrime. As these two effects work in the same direction, increased security will reduce the level of cybercrime (Van Eeten and Bauer 2008). However, it may trigger a new round in the technological race between cybercriminals and defenders.

3. Security incentives of stakeholders

The interdependence between stakeholders has, to a certain degree, contributed to a game of individual players blaming others for insufficient security efforts. Whereas Microsoft, due to its pervasive presence, is a frequent target (e.g., Schneier 2004; Perrow 2007), the phenomenon is broader: security-conscious ISPs blame rogue ISPs, ISPs blame software vendors, countries with higher security standards blame those with lacking law enforcement, and so forth. This is understandable but it does not represent an accurate picture of the complicated relations. A closer look at the incentives and the way externalities percolate through the system reveals a more nuanced picture. Several, admittedly imperfect, mechanisms, such as reputation effects, exist that better align private and social costs and benefits. Moreover, the high interconnectedness may result in the internalization of externalities at other stages of the value chain; the associated costs are therefore not externalized to society at large but to other players and may indirectly be charged back to the originators. For every player, multiple and sometimes conflicting incentives were at work, leaving the net effect contingent upon the relative influence of the specific drivers at work. Table 1 summarizes important incentives that shape the security decision of participants in the ICT value net.

Internet Service Providers (ISPs) are key market players and often criticized for poor security investment. Nonetheless, ISPs operate under several security-enhancing incentives although their effectiveness is mediated by the ISP's business model. Commercial ISPs will take security into account when it affects their revenue stream and bottom line. This is clearly illustrated by the example of viruses and spam. Early on, ISPs argued that emails were the personal property of recipients and that an inspection of the content of mails was a violation of privacy. Consequently, the responsibility for protecting their own machines and for dealing with spam was attributed to end users. With the exorbitant growth of spam the financial implications for ISPs also changed

fundamentally. Not only was the flood of spam a burden for their network infrastructure that would have required additional investment, the malware imported onto the network did indirectly affect the ISP's cost. Users of infected machines started to call the help desk or customer service at a fairly high cost per call for the ISP. Malicious traffic sent from infected machines triggers abuse notifications from other ISPs and requests to fix the problem, typically requiring even more expensive outgoing calls to customers. In extreme cases, the whole ISP could be blacklisted, causing potentially serious customer relations and reputation problems. Facing this altered economic reality, ISPs reversed their stance with little fanfare and started to filter incoming mail and to manage their customers' security more proactively.

Table 1: Security incentives of players in the ICT value chain

Player	Security-enhancing	Security-reducing
Internet Service Providers (ISPs)	Cost of customer support Cost of abuse management Cost of blacklisting Loss of reputation, brand damage Cost of infrastructure expansion Legal provisions requiring security	Cost of security measures Cost of customer acquisition Legal provisions that shield ISPs
Software vendors	Cost of vulnerability patching Loss of reputation, brand damage	Cost of software development and testing (time to market) Benefits of functionality Benefits of compatibility Benefits of user discretion Licensing agreements with hold-harmless clauses
E-commerce providers (banks)	Benefits of online transaction growth Trust in online transactions Loss of reputation, brand damage	Cost of security measures Benefits of usability of the service
Users	Awareness of security risks, realistic self-efficacy, exposure to cybercrime	Poor understanding of risks, overconfidence, cost of security products and services

In the present environment, ISPs operate under several more or less potent incentives. The strength of these incentives is mediated by the business model adopted by an ISP. "Rogue" ISPs whose business model is based on the hosting of shady activities will respond differently than commercial ISPs seeking legitimate business. Costs of customer support and abuse management as well as the cost of additional infrastructure requirement that might be required to handle floods of spam all have an immediate effect on the bottom line and will increase the incentives to undertake security-enhancing measures. Loss of reputation and brand damage work indirectly (and probably slower) but exert pressure in the same direction. ISPs are embedded in an interdependent system of service providers. If contacts via the abuse management desk are ineffective, other ISPs have a range of escalating options to retaliate for poor security practices with regard to outgoing malicious traffic, even if the origin are individual users.

Blacklists (or "blocklists"), inventories of IP addresses and email addresses reported to have sent spam and other forms of malicious code, are regularly used by ISPs to filter and block incoming traffic. Lists are maintained, typically by non-profit organizations such as SpamCop,

Spamhaus, or the Spam and Open Relay Blocking System (SORBS).[4] Other lists, such as those maintained by RFC-Ignorant detail IP networks that have not implemented certain requests for comment, for example, do not have a working "abuse@domain" address.[5] These lists are not uncontested and are sometimes seen as "vigilantes." But they are widely used and provide real-time countermeasures for network administrators. Each blacklist organization has procedures for delisting once a security leak is fixed. Substantial presence on one or more of these lists, reflected in the listing of many IP addresses belonging to an ISP for an extended time period, will drive up customer support and abuse management costs. It may also trigger reputation and revenue losses for an ISP. Both effects create an incentive to improve security measures, at least to respond in a timely fashion to abuse requests. In extreme cases, an entire ISP (and not just IP addresses or address ranges) may be blocked by blacklists and de-peered by other ISPs, raising the costs of this ISP significantly, possibly to the point where its business model becomes unsustainable.

All these incentives work in favor of enhanced security measures. On the other hand, the costs of increasing security, legal provisions that shield ISPs from legal liability, and the costs of customer acquisition all work in the opposite direction. Other things equal, they constitute incentives to adopt a lower level of information security. The net effect on ISPs is hence dependent on the relative strength of the components of this web of incentives. The high cost of customer calls (estimated by some ISPs to about $12 per incoming and $18 per outgoing calls), while providing an incentive to find alternative solutions to enhance security of end-users, also may provide an incentive to ignore individual cases that have not resulted in notifications to blacklisting services or abuse requests from other ISPs. Most ISPs estimated that only a few percent of the infected machines on their network show up in abuse notifications. Considerable advances in security technology, on the other hand, have enabled ISPs to move their intervention points closer to the edges of their network and thus automate many functions in a cost-effective way. In an escalating response, machines on the network may initially be quarantined with instructions to the user to undertake certain measures to fix a problem. Only in cases that cannot be solved in this fashion may customer calls become necessary. Overall, while the interdependencies among ISPs result in the internalization of some of the external effects, this internalization is imperfect. Hence, it is likely that the highly decentralized system of decision-making yields effort levels that fall short of the social optimum.

Software is a critical component in the ICT value chain. Exploits of vulnerabilities are one important attack vector. The market for exploits has become increasingly sophisticated as seen, for example, in the steady increase of zero-day exploits. Software vendors work in a complex set of potentially conflicting incentives whose overall net effect is difficult to determine. One strong factor working to enhance security efforts is the cost of vulnerability patching, which comprises the cost of patch development, testing, and deployment. As software is typically installed in many versions and contexts, the cost of patch development can be very significant. From the perspective of software vendors it is therefore often advantageous to invest in security upfront rather than have the follow-up costs of patching. However, given the complexity of modern software packages, it will not be economical to invest to the point where all potential flaws are eliminated. Loss of reputation and brand damage also strengthen the incentive to invest in security. However, reputation effects may work only slowly and are further weakened if the vendor has a strong market position, either because of a large installed base or because of limitations in the number of available substitute products. Nonetheless, as the security enhancements in the transition from Windows 2000 to Windows XP to Windows Vista illustrate, reputation effects do work.

[4] See http://www.spamcop.net, http://www.spamhaus.org/zen, and http://www.de.sorbs.net/.
[5] See http://www.rfc-ignorant.org/.

At the same time, there are several incentive mechanisms that, other things equal, weaken the effort to provide secure software. Time to market is lengthened by software testing and the cost of software development is increased. Thus, the more competitive the software market segment, the lower the incentive to improve security beyond a minimum threshold. Moreover, in software design complicated trade-offs have to be made between different design aspects. Security often is in conflict with functionality, compatibility, and user discretion (a benefit that many users covet). Lastly, hold harmless provisions in software licenses and shrink-wrap license agreements largely relieve software vendors from liability for financial damages stemming from software flaws. From the perspective of the whole value net, the diffusion of software for mobile devices and the expansion of open source software also increase vulnerabilities. The fear that increased liability of software vendors might reduce innovation rates is not fully unfounded but the strength of this effect is unknown.

Large businesses (firms with 250 and more employees) are a heterogeneous group. Many large business users have adopted risk assessment tools to make security decisions (Gordon and Loeb 2004). Their diligence will vary with size and possibly other factors such as the specific products and services provided. One particularly interesting industry is financial service providers. This is a rather diverse sector, encompassing different types of banks, credit card companies, mutual funds, insurance companies, and so forth. The rules for each of these players differ in detail. Focusing predominantly on merchant banks, Van Eeten and Bauer (2008) concluded that these financial service providers are to a considerable degree able to manage risks emanating from their customer relations. However, they need to make choices balancing enhanced security and the growth of their electronic business. In principle, they could use highly secure platforms to conduct ecommerce transactions. However, such an approach would likely have detrimental effects on users as it decreases the convenience of conducting business. Financial organizations thus face a trade-off between higher security and migrating transactions to cost-saving electronic platforms. Many financial service providers offer compensation for losses incurred by their customers from phishing or other fraudulent actions as part of this overall security decision. This practice aligns the incentives of the financial service provider with the goal of improved security (but not the incentives of individual users). Businesses other than financial service providers may often not be in a position to manage externalities associated with their clients. Therefore, more significant deviations between private incentives and social effects may exist, resulting in a sub-optimally low level of security investment by these firms.

Two other groups of players that deserve mentioning are small and medium enterprise (SMEs, typically defined as enterprises with fewer than 250 employees, including microenterprises) and residential users. Although this is a large and diverse group, these players also are in several respects similar. Like other participants, they work under multiple and potentially conflicting incentives. Unlike larger businesses that may be able to employ information security specialists, either in-house or via outsourced services, many SMEs and residential users have insufficient resources to prevent or respond to sophisticated types of attacks. Whereas awareness of security threats has increased, there is mounting information that many residential users underestimate their exposure and overestimate their efficacy in dealing with risks (LaRose, Rifon et al. 2005). Although these constitute similarities between SMEs and residential users there are also differences. In general, one can assume that businesses employ a more deliberate, instrumentally rational form of reasoning when making security decisions. In both cases, however, the benefits of security expenses will to a large degree flow to other users.

Individual businesses and users may suffer from the perception that their own risk exposure is low, especially if others protect their machines, the well-known free rider phenomenon. On the other hand, given increased information, a growing number of users in this category are aware of the threat exposed by breaches of information security. Thus, they realize to a certain extent that

they are the recipients of "incoming" externalities. Overall, one can expect that on average these classes of users will not be full free riders. Whereas some individuals and SMEs may over-invest, there is evidence that most will not invest in security at the level required by the social costs of information security breaches (Kunreuther and Heal 2003). This conclusion is corroborated by the observation that many individual users do not purchase security services, do not even use them when offered for free by an ISP or a software vendor,[6] and turn off their firewalls and virus scanners regularly if they slow down certain uses, such as gaming.

4. Analysis of the emerging patterns

Our findings indicate a number of market-based incentive mechanisms that contribute to enhanced security but also other instances in which decentralized actions are associated with externalities and hence may lead to sub-optimal. A pressing question is whether the response to malware of actors in information and communication markets is adequate or whether improvements are possible, either by private or public action (see Anderson, Böhme et al. 2008 for a compendium of possible measures). Pointing to a variety of reports that show increases in malicious attack trends, one might conclude that markets are not responding adequately. Our field work and analysis revealed a more diverse and graded picture.

Our findings suggest that all players work under some incentives that work in the correct direction. All market players we studied experienced 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. However, in many cases they are too weak or too localized to move their behavior towards more efficient social outcomes. In all cases, moreover, there are incentive mechanisms that undermine security. Overall, across the entire value net of all the different market players, three paradigmatic situations emerge:

(1) No externalities. This concerns instances in which a market player, be it an individual user or an organization, 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. Essentially, this scenario only applies to closed networks and user groups.

(2) 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 organization. As long as somebody in the value net internalizes 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 internalization. This scenario depicts a relatively frequent case and numerous examples were found where externalities were internalized by other market players. ISPs and financial service providers are two examples that were discussed in more detail above.

[6] For example, XS4All, a Dutch ISP, offered security services to their customers but less than 10 percent signed up. Eventually, the ISP decided to offer some of the protections by default. Many users do not use automatic updates to their software.

(3) 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, the other agents in the information and communication value net that absorb the cost are not in a position to influence them– 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 of ICT use may entail a significant opportunity cost for society at large if the delayed activities would have contributed to economic efficiency gains and accelerated growth.

The most poignant cases in this category are the externalities caused by lax security practices of end users. Some of these externalities are internalized 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 they themselves suffer consequences from the end user behavior, e.g., by facing the threat that a significant part of their network gets blacklisted. Estimates mentioned in the interviews suggested that the number of abuse notifications received by ISPs represents only a fraction of the overall number of infected machines in their network. This observation suggests that a considerable share of the externalities originating from ISP customers may not be mitigated.

Consequently, a large share of these costs of poor security practices of end users is borne by the sector as a whole and society at large, typically in the form of higher direct, indirect and implicit costs (see Van Eeten, Bauer et al. 2009). 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, in part due to insufficient information. While malware writers have purposefully chosen to minimize 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 personal information that can be used for financial gain. In that sense, end users should 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 is not perceived by the end user is an issue of incomplete information rather than a principal failure of the respective incentive mechanism.

The discussion so far has assumed that the threat level does not fully undermine the operations of the ICT value net. However, there is a risk of not just security failures or even a disaster but of catastrophic failure. Given the dependence of all aspects of global society on the Internet (and electronic communications in general), widespread and extended failure could certainly have catastrophic consequences. That such a pervasive failure or technological terrorism has not yet happened and has a low probability complicates the formulation of a response (Posner 2004). Like other events with a low but non-trivial probability, it could be considered a “black swan” event (Taleb 2007). Cost-benefit analysis of such catastrophic events would help in shaping more rational responses but it is extremely difficult. Complications include the choice of an appropriate time horizon, the quantification of the risk in question, problems of monetizing a wide range of qualitative impacts, and the determination of social discount rates applied to monetized future events. Nonetheless, assessing the potential costs of internet security breaches would greatly benefit from such an exercise.

5. Policy implications

The analysis presented in the preceding sections suggests that decentralized decision-making often generates correct incentives for stakeholders and that deviations from a desired security level often trigger the appropriate feedbacks. However, we also identified many instances where the resulting incentives were weak, too slow, or the net effects of the security-enhancing and security-reducing incentive mechanisms could be either positive or negative. Moreover, the entire ICT value net is under relentless attack from organized gangs and reckless individuals who dispose over increasingly powerful and sophisticated tools. Many of these activities are organized in countries where the cost of engaging in cybercrime are low: law enforcement is weak or non-existent, due to dire overall economic conditions, the opportunity costs of participating in cybercrime rather than pursuing other gainful employment are low, and technological means enable criminals to operate swiftly across many national boundaries. At the same time, numerous incremental measures to improve security beyond the status quo are known but may not be undertaken because they are afflicted with positive externalities. The benefits of such measures potentially help all stakeholders but their costs often need to be borne by one particular group. The entire ICT value net suffers from a prisoner's dilemma problem: everybody is worse off if decisions are made in a non-cooperative fashion. Where repeated interactions happen, it is partially overcome, as is illustrated by the cooperation among ISPs. Enhancing cybersecurity at a broader level will have to overcome this coordination and cooperation issue: it is a collective action problem. Whether such collective measures can be identified and successfully implemented without disadvantages that outweigh the potential improvements needs further examination.

There are two principal vectors for implementing such collective measures. They can target the realm of cybercrime or the information security decisions of stakeholders in the ICT value net. Both areas may be addressed with measures in four categories (or with a mix of such measures): legal and regulatory measures, economic means, technical solutions, and informational/behavioral approaches (see Table 2). Cybercrime can principally be reduced by increasing its costs and by reducing its benefits. Strengthening law enforcement via national legislation and multi-national and international treaties, such as the European Convention on Cybercrime, which has been ratified by 23 countries and signed by 23 others[7], is one important precondition to credible enforcement as it defines a legal basis for intervention. Equally important for the preventative effect of law enforcement are the forms and severity of punishment as well as the effectiveness and expediency of law enforcement. International collaboration is a particularly pesky obstacle that needs to be overcome in this area.

Given the fact that much criminal activity is organized in countries with relatively weak rules of law, other measures might be needed in addition to legislative and regulatory initiatives. A number of interesting proposals have been made to reduce spam via technical-economic mechanisms, such as the requirement to have a token or make a payment before a mailbox can be accessed (e.g., Loder, Van Alstyne et al. 2004). Although such measures are principally effective and interesting, they will only take care of malware disseminated via email. Moreover, their effectiveness will depend on a critical mass of users adopting the method. General measures to increase the opportunity costs of cybercrime, such as the creation of attractive employment opportunities for skilled programmers, will contribute to a reduction of activity long term. Many of the architects of the internet have proclaimed that it was not built for the current onslaught of legal and illegal activity. Several initiatives, such as DNSSEC (DNS Security Extensions) which

[7] See http://conventions.coe.int/Treaty/Commun/ChercheSig.asp?NT=185&CM=&DF=&CL=ENG (last visited February 28, 2009).

is designed to make the domain name system more secure, are underway. Several other protocols are deployed in the web and others may be rolled out as an underlay network. Whereas all of these measures already have triggered responses by attackers, they do make cybercrime more expensive and one would expect that this has a dampening effect, at least in the short term and other things equal. Lastly, measures of information sharing at a national level in CERTs and CSIRTs as well as the international level, such as in the more than 200 organizations collaborating in FIRST (Forum for Incident Response Teams) are important steps in the right direction.

Table 2: Principal policy instruments to enhance information security

Predominant policy vector	Cybercrime	Information security
Legal and regulatory measures	• National legislation • Bi-and multilateral treaties • Forms and severity of punishment • Law enforcement	• National legislation/regulation of information security • Legislation/regulation of best practices to enhance information security • Liability in case of failure to meet required standards • Tax credits and subsidies
Economic	• Measures that increase the direct costs of committing fraud and crime • Measures that increase the opportunity costs of committing fraud and crime • Measures that reduce the benefits of crime	• Level of financial penalties for violations of legal/regulatory provisions (compensatory, punitive) • Payments for access to valuable information • Markets for vulnerabilities • Insurance markets
Technical	• Redesign of physical and logical internet infrastructure	• Information security standards • Mandated security testing • Peer-based information security
Informational and behavioral measures	• National and international information sharing on cybercrime	• National and international information sharing on information security • Educational measures

Whereas measures devised to fight cybercrime are relatively uncontested, this is not true for policies directed toward information security markets. Nonetheless, there is an emerging discussion in many countries as to whether such measures are appropriate, whether they should be targeted a key players in the value net, and which organizations should be in charge of such initiatives (Anderson, Böhme et al. 2008; OECD 2009). Several countries have adopted laws against spam. Although their effectiveness is sometimes questioned, legislation such as the US CANSPAM Act of 2003 and subsequent implementation measures by the Federal Communications Commission (FCC) and the Federal Trade Commission (FTC) have provided a legal basis to prosecute several spammers with deterrent effects. The recent discussion goes beyond such measures to explore specific regulatory intervention. A critical weakness of any attempt to legislate or regulate security is that specific measures will be outsmarted by new attack

technologies very quickly. However, the notion of "best practices" could be developed on an ongoing, adaptive basis. This is, for example, the approach taken by the Australian Communication and Media Authority (ACMA) in the Australian Internet Security Initiative. In this collaborative effort, 59 ISPs (as of February 28, 2009) and the regulatory agency share information to reduce the threat from botnets. The approach is not without critics, many of whom point to the lack of transparency in following up on security threat reports. Nonetheless, the ability of the regulatory authority to threaten with the imposition of formal regulation seems to have boosted participation and the initiative is expanding steadily.

An even more contested issue is the modification of existing liability rules. In principle, liability rules would facilitate a desirable evolutionary learning process. By creating enforceable rights and obligations, they would allow common and case law institutions to gradually develop a body of best practices and reshape incentives to reduce negative and strengthen positive externalities. Compared to the status quo, such rules will create advantages for some players and disadvantages for others. Most likely, such changes will activate veto players of losers that will attempt to block solutions even if the winners could have compensated them. Changes in economic institutions and incentives could achieve similar outcomes. Among the possible measures discussed in the literature are increased penalties for security flaws, the establishment of markets for vulnerabilities, and insurance markets. All these proposals have some appealing features but most also have serious downsides (see the more detailed discussion in Van Eeten and Bauer 2008). At a technical level, information security standards, mandated security testing, or peer-based information security approaches could be considered. Most nations are actively engaged in information sharing programs at the national level. Some have established reporting requirements for security breaches. Early observations suggest that such measures work in favor of enhanced security. At least, they increase market transparency for consumers.

The effectiveness of the available broad spectrum of measures is not well established. Very few empirical studies document the possible impact of specific measures. As it is unlikely that a comprehensive information base that will allow such judgment is available any time soon, measures might have to be adopted on an experimental, trial and error basis. Different options may complement each other (such as national and international legislation) whereas others may be partial or full substitutes for each other. More research and practical policy experimentation is needed to move this process along.

6. Conclusions

Information and communication technologies form a vast, in many areas relatively open ecosystem. This openness is, to a certain degree, a main ingredient into the innovative dynamics and success of the internet. At the same time, it renders the infrastructure and associated services vulnerable to attack. The realms of cybercrime and information security can be modeled as two distinct but co-evolving markets. Based on extensive interviews with experts from the private, non-profit, and public sector, this paper develops a compact view of the incentives of individual stakeholders and categories of stakeholders to undertake measures that enhance security. All the participants in the ICT value chain operate under security-enhancing and security-reducing incentives. The overall system has many feedbacks that enhance security but externalities, uncompensated interdependencies between players are also pervasive. Three typical cases describe the extent and severity of externalities. In some cases, all social costs and benefits are internalized and no externalities are present. In other cases, one stakeholder in the value net other than the source of the externality is able to control it, at least to a certain degree. In yet other cases, no such opportunities exist and externalities are imposed on the whole value net and society at large. A vast spectrum of possible policy measures is available to combat cybercrime

and support information security. Too little is known about the effectiveness of these instruments, forcing policy-makers and private sector players to proceed with multiple means at once in a trial and error fashion. From an evolutionary perspective, this is a rational strategy (comparable to a broadband antibiotic) as policy diversity may allow learning as to which approaches or which mixed of approaches are most effective. Systematic learning requires that the effects of measures are observed on an ongoing basis to generate deeper understanding as to the relative effects of different solutions. Most likely, an effective overall approach to better realigning the economic incentives of players in the ICT value net will have to combine measures at the legal, economic, technical and informational level.

References

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.

Anderson, R., R. Böhme, et al. (2008). Security Economics and European Policy. Cambridge, University of Cambridge Computer Laboratory.

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

Bauer, J. M., M. Van Eeten, et al. (2008). Financial Implications of Network Security: Malware and Spam, Report for the International Telecommunication Union (ITU), Geneva, Switzerland, available at http://www.docstoc.com/docs/2399824/ITU-Study-on-the-Financial-Aspects-of-Network-Security-Malware.

CSI (2008). 2008 CSI Computer Crime and Security Survey. San Francisco, CA, Computer Security Institute.

Gordon, L. A. and M. P. Loeb (2004). The Economics of Information Security Investment. Economics of Information Security. L. J. Camp and S. Lewis. Dordrecht, Kluwer Academic Publishers**:** 105-128.

Illing, G. and M. Peitz, Eds. (2006). Industrial Organization and the Digital Economy Cambridge, MA, MIT Press.

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

Kanich, C., C. Kreibnich, et al. (2008). Spamalytics: An Empirical Analysis of Spam Marketing Conversion. Proceedings of the 15th ACM Conference on Computer and Communications Security, Alexandria, VA.

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

LaRose, R., N. Rifon, et al. (2005). Understanding Online Safety Behavior: A Multivariate Model. International Communication Association New York.

Loder, T., M. Van Alstyne, et al. (2004). An Economic Answer to Unsolicited Communication. Proceedings of the 5th ACM Conference on Electronic Commerce New York, NY.

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

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

Posner, R. A. (2004). Catastrophe: Risk and Response. New York, NY, Oxford University Press.

Schipka, M. (2007). The Online Shadow Economy: A Billion Dollar Market for Malware Writers, White Paper, MessageLabs Ltd.

Schneier, B. (2004). Secrets and Lies: Digital Security in a Networked Society. New York, Wiley.

Shapiro, C. and H. R. Varian (1999). Information Rules: A Strategic Guide to the Network Economy. Boston, MA, Harvard Business School Press.

Taleb, N. N. (2007). The Black Swan: The Impact of the Highly Improbable, Random House.

Van Eeten, M. and J. M. Bauer (2008). The Economics of Malware: Security Decisions, Incentives and Externalities, Directorate for Science, Technology and Industry, Committee for Information, Computer and Communications Policy, DSTI/ICCP/REG(2007)27, Paris: OECD, available at http://www.oecd.org/dataoecd/53/17/40722462.pdf.

Van Eeten, M., J. M. Bauer, et al. (2009). Damages from Internet Security Incidents: A Framework and Toolkit for Assessing the Economic Costs of Security Breaches, Delft University of Technology, The Netherlands.

Varian, H., J. Farrell, et al. (2005). The Economics of Information Technology: An Introduction. Cambridge, MA, Cambridge University Press.