

THE ECONOMICS OF MALWARE

Michel J. G. van Eeten*
Johannes M. Bauer**
John P. M. Groenewegen*
Wolter Lemstra*

Delft University of Technology*
Michigan State University**

Correspondence to:
M.J.G. van Eeten
Faculty of Technology, Policy and Management
Delft University of Technology, PO Box 5015, 2600 GA Delft, The Netherlands
m.j.g.vaneeten@tudelft.nl, T: +31 (0)15 2787050, F: +31 (0)15 2786233

August 17, 2007

Prepared for presentation at the
35th Telecommunications Policy Research Conference
Arlington, CA, September 28-30, 2007

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 around combating cybercrime, developing Computer Emergency Response Teams (CERTs), raising awareness, information sharing, fostering education and other initiatives.

During the same period, security threats have increasingly captured the public's attention – fueled 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 a 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 and malicious mobile code.

The effects of malware have exploded in the last few years, forcing us to rethink the way in which information security is pursued. For governments, increasing public attention implies increasing political pressure to intensify their actions, beyond the initiatives already underway. The question is: How? What policies and initiatives are needed?

How to improve cybersecurity is hardly a straightforward question. Notwithstanding rapidly growing investments in security measures, it has become clear that cybersecurity is a technological arms race that will not be decided in the immediate future. 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 the 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, to say the least. Furthermore, security measures may also impede innovation and productivity. Those involved in cybersecurity frequently tend to overlook that the reason why the internet is so susceptible to security threats is the same 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 very start, credit card companies have struggled with rising fraud. That hasn't stopped them from expanding their online business. The benefits of that growth were consistently higher than the costs of the increase in fraud that came with it. Rather than implementing far-reaching security measures that would restrict the usefulness of the

system, they've adopted strategies to fight instances of fraud, up until the point where the costs of further reductions in fraud are higher than the remaining damages.

All this means that total security is neither achievable nor desirable. Actors need to make their own tradeoffs regarding what kind of security measures they deem appropriate and rational, given their business model. Clearly, these business models are very different 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 organizations to end users and beyond.

In other words, many instances of what could be conceived as security failures are in fact the outcome of rational economic decisions, given the costs and benefits facing the actors involved. What is needed, then, is a better understanding of these costs and benefits – in short: of the economics of cybersecurity. This report outlines a research project to this aim, preparing OECD members for the next generation of policies, as well as providing a better foundation for the public-private partnerships set up to deal with cybersecurity.

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 a focus on technological causes and solutions.

Most of this innovative research has yet to find its way into the realm of policymakers, let alone into the policies themselves. While reports like the OECD survey on the culture of security (OECD, 2005) generally recognize that there is more to cybersecurity than technology, the proposed measures are still mostly oriented in that direction: developing technological responses and efforts to stimulate their adoption. Think of initiatives to promote authentication, encryption and Trusted Third Parties, awareness campaigns urging people to improve the security of their systems, certification schemes tied to security standards, and clearinghouses for information on security threats and their remedies such as CERTs.

Notwithstanding the necessity of these initiatives, they typically ignore the economics of 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, the incentives that stimulate efficient behavior are 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 nineties, when the scale of virus distribution was rapidly increasing and many end users (home, corporate, governmental) were affected, most ISPs argued that virus protection was the responsibility of the end users themselves. The computer was their property, after all. They further argued that they couldn't scan the traffic coming through their e-mail servers, because that would invade the privacy of the end user. The mail message was also considered the property of the end user. About five years ago, this started to change. The spread of viruses and worms had grown exponentially and now the infrastructures of the ISPs themselves were succumbing to the load. 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 that they identified as malignant. The effects of the property rights had been extended: the property rights of the infrastructure now gave the incentive to invest in fighting malware. One could view this as an example of an invisible hand: self-interested behavior of ISPs led to a more thorough defense against email-based viruses and increasing net social benefits.

In many cases, an economic perspective on cybersecurity – and malware in particular – provides us with more powerful analysis and a 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. More work is needed, however. As we will see, most of the research so far has been based on the methods of neoclassical 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.

We discuss the implications of these neoclassical and new institutional approaches in more detail in the next chapter. For now, we briefly key mention three limitations: (1) they provide limited insight into how actors actually perceive the cost, benefits and incentives they face; (2) they have difficulties taking into account dynamic and learning effects, such as how a loss of reputation changes the incentives an actor experiences; and (3) they treat issues of institutional design as somewhat trivial. That is to say, the literature assumes that its models can indicate what market design is optimal, that this design brought into existence at will and that actors will behave as the model predicts. If the past decade of economic reforms –such as privatization, liberalization and deregulation – have taught us anything, it is that designing markets is highly complicated and sensitive to context. It cannot be based on formal theoretical models alone. Institutional design requires an in-depth empirical understanding of current institutional structures.

To provide the basis for new policies, we propose to complement the state-of-the-art understanding of the economics of malware with qualitative field research that provides empirical evidence on the way in which actors actually make security tradeoffs, how they perceive their institutional environment, the incentives they face and how these have changed, as well as the externalities that arise from these incentive structures.

The remainder of this chapter locates our project within the context of current security and privacy guidelines, as well as two adjacent economic issues that are outside the scope of this project: the criminal business models underlying malware and the overall economic impact of malware. Chapter II first presents the state of the art of the economics of cybersecurity, with a focus on incentives and externalities. We then review the underlying assumptions and the gaps they imply for our current understanding of the economics of malware.

II. Economics of information security and the OECD Guidelines

In 2002, the OECD released their “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 network, 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”. According to the OECD, the guidelines reflect the shared understanding of member countries as well as a variety of business and consumer organizations.

The “culture of security” that the guidelines aim to promote will be influenced by the incentive structures surrounding security tradeoffs. Yes, the focus on security may 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. Key points include (OECD 2006b):

- 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.

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 – that in turn may reduce their own investment below the optimal level. 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.” But many software markets do not reward such behavior.

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 ensuring the security of the software. 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, and instead uses the host to attack other systems.

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 – that in turn may reduce their own investment below the optimal level. 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.” But many software markets do not reward such behavior. 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 ensuring the security of the software. 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, and instead uses the host to attack other systems.

In short: the so-called “culture of security” is very sensitive to 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.

III. Related economic issues

Our project approaches malware from an economic perspective, with a focus on the incentives and externalities that influence the security tradeoffs of market players. This focus means that several relevant economic issues are outside the scope of the project. Two of these issues are the overall economic impact of malware and the business models that drive the production of malware. We briefly discuss them here, as they shed some light on context in which our project is located.

Economic impact of malware

Currently, there are no authoritative data on the overall impact of malware. The annual *Computer Crime and Security Survey* by the Computer Security Institute and the FBI is widely regarded as

the best data available on the damage of breaches in cybersecurity at organizations (Gordon et al. 2006). The survey has found that virus attacks are the leading cause of financial losses. The organizations that were willing to (anonymously) share their estimates, reported an average loss per firm of \$167,713 in 2005 – which adds up to a total of \$52,494,290 for the 313 respondents that were willing and able to estimate losses.

The reliability of other data is more controversial, but typically the figures are in the range of billions of dollars. A study released in July 2000 by InformationWeek Research and PricewaterhouseCoopers estimated that the cost of malware exceeded US \$1.5 trillion worldwide that year. The impact on US businesses with more than 1000 employees is estimated to be \$266 billion or approximately 2.7% of the Gross Domestic Product (Cavusoglu et al. 2004b).

In the same year, there were Distributed Denial of Service (DDoS) attacks against sites like Amazon, Yahoo and eBay. These DDoS attacks were made possible by botnets – networks of PCs comprised by malware and under the control of an attacker. One analyst assessed the damage of these attacks at \$1.2 billion. About \$1 billion of that amount was the result of a negative impact on stock prices, about \$100 million was lost revenue from sales and advertising, and about \$100 to \$200 million were put into security upgrades (Nicolai 2000). However, the victims saw it rather differently. Yahoo, for example, argued that it did not suffer major losses and that the losses in advertising could in part be recovered by replacing their own ads with those of paid clients (Denning 2000).

Weaver and Paxton (2004) have attempted to estimate how much damage an attacker with extensive resources – such as a nation state – could do to the United States with a worm-based internet attack. Adding up lost productivity, repair time, lost data, and damage to systems, they argue that “it is not implausible to conceive of attacks that could disrupt 50 million or more business computers, causing tens or perhaps hundreds of billions of dollars in direct damages.” These numbers do not include possible indirect effects on other infrastructures – such as the reported cases of worms affecting an ATM-network and the safety systems of a nuclear power plant (Poulsen 2003a; Poulsen 2003b). While Weaver and Paxton set out to develop a plausible worst-case scenario, security firm Mi2g has produced similar figures, but then claiming that this was the actual damage caused by a variety of worms in 2003 (Thompson 2004).

It is important to note, however, that all of these figures are controversial. Sometimes the victims themselves may disagree. While they may have incentives to downplay the damage, Anderson (2002) has argued that security experts have incentives to exaggerate the estimates. Typically, they are used to create a sense of urgency – calling for more security investments or some form of governmental intervention. The problem with all these estimates – no matter how reliable – is that they remove the damage assessments from the context in which they matter most: decentralized cost/benefit tradeoffs by actors, corporate or otherwise. One can put a monetary value on the overall economic impact of a certain security failure, but that number really is not very informative, unless placed in a context where it is connected to estimates of what it would cost to prevent that damage. This requires a focus on the decentralized tradeoffs of market players. Only that way does it become clear whether alarming damage estimates indeed imply that investments need to be stepped up.

Malware business models

The production of malware is increasingly a profit-driven enterprise. In fact, profit has been a key driver behind the exponential growth of malware in recent years, according to many experts. This has raised interest in how malware actually makes money – in other words, what its underlying business model is. Much of this business model is criminal in nature, though the money trail sometimes unwittingly overlaps with those of bonafide market players. Recently, Microsoft was forced to apologize for the fact that one of its own advertisement channels has displayed a banner ad for a software application called “Winfixer,” which Microsoft itself said was malware (Kirk 2007).

The fact that most of the business model is criminal in nature makes it rather difficult to provide reliable numbers. Also, most malware can be used for a variety of criminal purposes – i.e., business models. Botnets, for example, are nothing more than tools. An analysis of the HoneyNet Project found a wide variety of uses (Bächer et al. 2005):

- Distributed Denial-of-Service Attacks. Often botnets are used for DDoS attacks on a computer system or network. Such attacks cause a loss of service to users, typically the loss of network connectivity and services by consuming the bandwidth of the victim network or overloading the computational resources of the victim system. This has been used, among other purposes, to extort money from victims (Sturgeon 2005) and to disrupt the services of competitors (Department of Justice 2004).
- Spamming. Some botnets turn the infected hosts into proxies which can be used to send massive amounts of spam. Some bots also implement a special function to harvest email-addresses. It is estimated that currently, the bulk of the total volume of spam originates from a botnet (Jan 2007).
- Sniffing Traffic. Botnets may also install packet sniffer at the infected hosts. The sniffers intercept interesting clear-text data passing by, mostly aiming to retrieve sensitive information like usernames and passwords.
- Keylogging. Beyond sniffing the network packets on the infected computer, some botnets employ keyloggers. These pieces of malware capture the actual key strokes of the user of the compromised PC. Harvesting information from key strokes can be supported by filtering mechanisms – e.g., a filter that looks for key strokes in the vicinity of the keyword ‘paypal.com.’
- Spreading new malware. In most cases, botnets are used to spread new bots through a variety of tactics.
- Installing Advertisement Addons and Browser Helper Objects. Some advertisers pay the websites that host their ads for every instance someone clicks on them. Attackers can strike a deal with the hosting website to instruct the bots in the botnet to automatically click on the advertisements, generating thousands of clicks in an instant. This process can be further enhanced if the bot hijacks the start-page of a compromised machine so that the “clicks” are executed each time the victim uses the browser.
- Google AdSense abuse. A similar abuse is also possible with Google’s AdSense program. The botnet is used to artificially generate clicks on Google advertisements, thereby generating revenue for the hosting website.

- Attacking IRC Chat Networks. Botnets are also used to attack Internet Relay Chat (IRC) networks by flooding the network with service requests. In this way, the victim is brought down in a way similar to a DDoS attack.
- Manipulating online polls/games. Every bot has a distinct IP address, so every vote will have the same credibility as a vote cast by a real person. This can be used to manipulate the outcomes of polls and games.
- Mass identity theft. Often the combination of different functionality described above can be used for large scale identity theft: the botnet can send out phishing mails, host fake website that mimic legitimate sites, install keyloggers and sniffers – all to capture private information to enable identity theft.

The exact business models underlying these kinds of attacks are difficult to piece together. Occasionally, interesting snippets of information become available that shed some light on the economics of these forms of crime. In 2005, a company based in Russia used an existing legitimate business model – an affiliate style program – to exploit vulnerabilities. The company was paying participating Web sites 6 cents for each machine they infect with adware and spyware. To participate, the website were asked to place an exploit on their sites. A security experts estimated that the company could collect as much as \$75,000 annually from the adware it placed on the infected machines, which cost it about \$12,000 in payments (Keizer 2005).

Security firm TrendMicro reported that experts who have monitored IRC chat rooms found that a DDoS attack can cost between \$500 and \$1,500, while smaller botnet attacks are priced between \$1 and \$40 per compromised PC (Trendmicro 2007). In 2006, the FBI arrested a ‘herder’ who rented out his botnet, complete with guidelines on how many bots would be needed to crash corporate webs of various sizes. The minimum rate was 10,000 bots at four cents a piece – i.e., \$400 (FBI 2006).

The economics of spam have been more amenable to research. Many experts have calculated that spam can be very profitable even at almost negligible response rates. Mailchannels, a mail technology provider, estimated that if 100,000 messages are needed to generate \$1 of revenue, then a botnet of 10,000 PCs is sufficient to earn about \$4,000 per day (MailChannels 2007). Recently, a Dutch spammer was fined €75,000 for sending spam through a botnet (Leyden 2007). The botnet consisted of 600-700 compromised PCs, which over course of 14 months had sent out about nine billion spam messages, reportedly earning him an estimated € 40,000 before he was arrested.

New variants constantly find different ways to make money. In 2006, we have witnessed a surge of stock spam. These so-called pump-and-dump schemes use a flood of spam messages to raise interest in certain thinly traded stocks, hoping that this will trigger price hikes. Once the price hike has occurred, they dump the stocks. Böhme and Holz (2006) have analyzed these schemes and found that the spam touting a certain stock is followed by increased trading activity of the cited stock and positive cumulative abnormal returns. Frieder and Zittrain (2007) also conclude: the pump-and-dump business model is indeed successful.

The reports may be scattered, but many agree that the implication is clear: Malware has become the point of convergence of computer hacking and organized crime. Some even estimate that the revenue from malware is eclipsing the revenue of anti-virus vendors (Leyden 2006). The growing

involvement of organized crime professionalizes malware production. Sites offering customized malware – so-called Trojan supermarkets – allow criminals to order malware tailored to their purposes. All of this will undoubtedly exacerbate the threat that malware poses to bonafide market players, as well as the magnitude of the externalities that market players may impose on each other while responding to this threat.

IV. Economics of information security: reflecting on the state of the art

In 2001, Ross Anderson published a path-breaking paper entitled “Why Information Security is Hard: An Economic Perspective” (Anderson 2001). It identified different research initiatives that showed why economic methods were exceptionally powerful to explain and address issues of information security. A handful of other scholars had come independently to the same realization. In the following six years, over 200 papers have appeared, turning the fragmented efforts of a few scholars into a burgeoning research community. Much of this work has been brought together through the yearly Workshop on the Economics of Information Security.

We have surveyed the existing literature on the economics of information security, with special emphasis on research that uses the concepts of incentives and externalities. This chapter first summarizes the state of the art by describing the typical security problems that are addressed, the way in which incentives are analyzed and the types of externalities that are identified. Next, we identify the literature’s underlying assumptions, typical outcomes and proposed solutions. We conclude by discussing these findings and gaps in current research on incentives and externalities. All of this builds toward the development of an empirical investigation into the economics of malware – which is the topic of the next chapter.

A wide variety of security issues are tackled from an economics perspective. Much of the work, however, clusters around six topics which all involve potential externalities:

- security of end users
- security of the firm
- security of the network
- security of software
- privacy
- digital rights management

All of these clusters are relevant to the issue of malware, albeit in different ways. Combined, they highlight the multi-faceted nature of malware as well as the underlying economic mechanisms that may hinder or encourage its reduction.

Security of end users

Much has been written on the rise of botnets – networks of compromised end user PCs, in some cases spanning millions of systems, which can be controlled by an attacker. Many end users do

not adequately protect their system, making them vulnerable to a variety of malware. This behavior can be explained by different factors, such as a lack of awareness or technical competence and user habits that are not adapted to the new networked environment. But user incentives are at least as important. Many of the negative impacts of inadequate end user security are suffered by other actors than the end users themselves (Varian 2000; Anderson and Moore 2006). A typical payload may be a software to initiate denial of service attacks against, for example, Microsoft or online retailers. Many instances of sophisticated malware actively try to minimize their impact on the infected host, so as not to trigger detection and removal. As a result, incentives are misaligned and externalities persist. End users bear the costs of improving the security of their systems, such as a subscription to antivirus software, while others actors reap many of the benefits.

These issues affect not only to home users but also, for example, server administrators. Their awareness and competence is arguably higher, but still patches for software vulnerabilities are not applied in time, even when they are available well before malware actually emerged to exploit the vulnerability. August and Tunca (2006) use cost to explain why even server administrators do not maintain proper patching procedures. Patching is time consuming and therefore costly – some estimates run into the hundreds of dollars per patch per server. They explore different mechanisms that might provide the incentives to mitigate insecure patching practices. In addition, Png, Tang et al. (2006) argue that the inertia among end-users in taking precautions, even in the face of grave potential consequences, is explained by the fact that patching efforts are strategic substitutes: the higher the patching effort of others, the lower the patching effort of any particular individual, an incentive structure contributing to free rider problems.

Others aspects of end user behavior have also been researched through this lens. For example, Camp and Lewis (2004, p. 197) argue that often it is rational for users to subvert security measures. When information security means ensuring that end users have no place to hide their own information, or when security is implemented to exert detailed control over employees, then it provides users with perverse incentives which move them to resist this control.

Security of organizations

A large cluster of research has focused on determining the optimal security investment levels for organization, mostly notably firms. Not all insecurity is worth preventing. Anderson (2002) questions the widely held idea that there is underinvestment in security. Drawing on parallels with environmental economics, he argues that security community has built-in incentives to overstate the problem. Anderson concludes that many firms get it about right. For example, preliminary analysis for software vulnerabilities suggest that the so-called “return on security investment” rates are around 12-21 percent – which is hardly a sign of massive underinvestment (Soo Hoo et al. 2001). A similar finding is reported by Choi, Fershtman, and Gandal (2005), who find that the instances in which firms make sub-optimal decisions are limited.

There have been numerous attempts to assess the costs of security breaches. Several projects have studied effects of disclosing security breaches on the capital market value of those corporations. A large portion of cybersecurity breaches does not have a significant economic impact on organizations. Several studies have looked at the effect of security breaches on stock market

values of corporations. Campbell et al. (2003) found that, on average, only breaches of confidentiality had a significant negative impact. In other cases, the effect was not significant. For these cases, the average decline of market value was about 5 per cent. Cavusoglu et al. (2004a) are more concerned, concluding that the cost of poor security is very high for investors. Their analysis found 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 percent of their market value within two days of the announcement—an average loss in market capitalization of \$1.65 billion per breach. Security developers, on the other hand, gain about 1.4 percent in market value after an announcement.

In light of scarce resources and difficult tradeoffs, more formal methods have been developed to support management decisions on how much to invest (Tiwari and Karlapalem 2005). Gordon and Loeb (2006) argue that, contrary to popular belief, these decisions lend themselves to cost-benefit analysis and that the optimum level of cybersecurity investment is where the marginal costs of increased information security equal the marginal decrease in the costs due to events such as worm and virus attacks, hacking and confidentiality breaches. A number of financial tools, such as Return on Investment, Net Present Value, Internal Rate of Return and Annual Loss Expectancy, are available to manage these tradeoffs. The idea is to rationalize investment decisions, whereas many firms now are believed to be driven by qualitative information, past incidents, rules of thumb or other approaches.

Recently, this research is being complemented by empirical field work at firms, to determine how firms actually make these kinds of decisions. Dynes et al. (2005) found that the managers in the firms they studied believe that information security is less a competitive advantage than a qualifier for doing business. The main drivers for adopting additional information security were customer requirements and government regulation – a finding that was confirmed by Rowe et al. (2006), who studied a wider variety of organizations. None of the interviewed firms felt that a lack of information security on their part would result in their being liable for damages, with the possible exception of liability resulting from Sarbanes-Oxley.

Security of networks

Closely related to the previous cluster, there are research efforts to model the security investments within networks of organizations. If security of an organization is partly dependent on the actions of other organizations within the network, then this changes their incentive structure for investments. There are a substantial number of studies that use game theoretic models to analyze the effects of these incentive structures and the resulting externalities. According to Varian (2004), if system reliability depends on the effort of many individuals, that makes it a public good. It is well-known that the provision of public goods may result in free rider problems, as individuals tend to shirk. Garcia and Horowitz (forthcoming) have built a model for ISPs, which predicted that there is a serious potential for underinvestment. When risk is interdependent, Kunreuther and Heal (2003) argue, then security investments can be strategic complements: An individual taking protective measures creates positive externalities for others that in turn may reduce their own investment below the optimal level.

Researchers have explored alternative institutional arrangements that provide incentive structures which reduce these externalities, such as insurance or information sharing. Schneier (2004) has argued that insurance is an essential part of internalizing security externalities. First, he says, we need to assign and enforce liabilities. Next, allow parties to transfer liabilities. This will create demand for insurance. He predicts that insurance in turn provides organizations with incentives to raise their security investments and creates demand for secure products, as insurers will charge different premiums for different levels of security. Böhme (2005) has looked at actual insurance policies for cybersecurity and comes to a similar conclusion as Schneier. Insurance gives firms an incentive to increase their level of security – not only because of differentiating premiums according to different classes of risk, but also because insurance companies have an incentive to reinvest a fraction of their revenues to improve their base of information, which finally yields new insights for more secure products. That said, there are still numerous problems from realizing this potential. The market is currently underdeveloped and underused.

Game theoretic models have been used to study the effects of the strategy of information sharing – i.e., organizations exchanging information related to computer security breaches, as well as to failed breach attempts. Information regarding the methods for preventing, detecting and correcting security breaches is also presumed desirable because it helps organizations to learn from the breaches experienced or prevented by other organizations. The modeling effort of Gordon et al. (2003) has shown that when security information is shared among firms, each firm reduces the amount spent on information security activities. Nevertheless, it can still lead to an increased level of information security, because the shared information allows firm to reach that level at a lesser cost.

Security of software

The security of software, or the lack thereof, has been the subject of a heated debate. As malware and many security threats prey upon software vulnerabilities, researchers have been studying the incentives under which software is produced. It is generally agreed that software manufacturers have insufficient incentives to produce secure software (Anderson and Moore 2006). This has led many to view vulnerabilities as negative externalities, similar to environmental pollution. The market rewards 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 ensuring the security of the software. Also there are positive network effects when the software is used by more people – more consumers and more developers of third-party applications. Stringent security practices could make life more difficult, especially for the latter group.

One area of research in this cluster has focused on vulnerability disclosure: How should benign users disclose discovered vulnerabilities to achieve optimal social outcomes? Disclosing it to the vendor only has the disadvantage that vendors sometimes respond slowly or not at all and the vulnerabilities remains unfixed for a long time. Disclosing it to the public gives the vendors a strong incentive to immediately fix it, but also provides hackers with a window to exploit the vulnerability. Arora et al. (2004) have collected empirical data which demonstrates that public disclosure made vendors respond with fixes more quickly, while at the same time the number of attacks increased. However, they also found that the number of reported vulnerabilities did decline over time. The latter point is still undecided. Is the pool of vulnerabilities depleted over

time as vulnerabilities are discovered? Ozment and Schechter (2006) argue that, in the case of OpenBSD, a security focused operating system, vulnerabilities did indeed decrease over time. In other words, security improves with the age of the software – leading them to conclude that it resembles wine rather than milk. Rescorla (2004), on the other hand, could not confirm that finding, leading him to question the usefulness of finding and disclosing vulnerabilities. Current work on disclosure models focus on hybrid practices, such as those of CERT, which notify the vendor and after a grace period, make the vulnerability public (Cavusoglu et al. 2005; Nizovtsev and Thursby 2005). Another area of research has focused on an alternative approach: designing a market for vulnerabilities. Researchers have studied different designs for such a market (e.g., Camp and Wolfram 2004; Ozment 2004; Schechter 2004).

V. Externalities

Externalities are at the heart of the incentive issues discussed in the previous sections and are intricately related to information security problems. Most fundamentally, externalities are interdependencies between agents that are not reflected in market transactions (payments, compensation). A positive externality exists, if an activity by agent A (such as an investment into improved information security) not only creates benefits for A but also for other agents B. Conversely, a negative externality exists if an activity by agent A (such as avoiding an investment into improved information security) not only creates costs of security violations for A but also for other agents B. In the formulation of the mainstream economic model, these interdependencies lead to deviations from an optimal allocation of resources. Negative externalities result in an overuse of a resource or overproduction of a good or service compared to the social optimum whereas positive externalities lead to an underuse or underproduction of the resource afflicted with the externality (Friedman 2002, p. 599). Which phenomena constitute externalities depends on the specification of legal rights and obligations in the status quo. Therefore, one way to overcome their undesired effects is to modify the legal and institutional framework of individual decisions (to “internalize” externalities).

In the public policy literature, 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, p. 527). An alternative typology distinguishes between technological and pecuniary externalities (Bobzin 2006). 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) by other agents. Early contributions to the subject, for example, by Marshall or Pigou during the late nineteenth and early twentieth century, treated externalities as an exception, a rare 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 networked computer environments, which raise several new and unique issues. The high degree of physical and logical 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 or other services. Many users of cyberspace view it as a commons, in which transactions take place according to gift logic rather than marketplace logic which creates an inherent bias against market-based solutions. Moreover, increasing information security is further complicated by asymmetric information issues. Often, for example in the case of Trojan horses or botnets, externalities are generated without the explicit consent nor knowledge of an individual user. All these factors influence the prevalence of externalities and possible ways to address them.

Origins of externalities in networked computer environments

External effects may originate at different stages of the value chain in networked computer environments. Depending on the source of the externality, the individual decision-making calculus causing the externality may be slightly different. It is common to assume that economic agents focus on their own costs and benefits and largely neglect costs or benefits of third parties.¹ Table 1 provides an overview of the sources of externalities in networked computer environments. The table captures the main, but not necessarily all stakeholders. Agents in the column are the sources of externalities whereas agents in the rows are the recipients. Not all agents 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 cause an externality for another agent in the same category. For example, the lax security policy of one ISP may cause externalities for other ISPs.

A first source of possible externalities is software vendors. Several authors have pointed out that, when deciding the level of investment in activities that reduce vulnerabilities, software vendors will primarily take their private costs and benefits into account (e.g., Schneier 2004). This does not necessarily imply that software developers will ignore costs imposed upon other stakeholders caused by flawed code. 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 users of the software. However, it is not self-evident that such reputation effects are sufficient to fully internalize externalities. The situation is aggravated by the unique economics of information markets, most importantly their high fixed 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; Anderson 2002; Shostack 2005) to take advantage of first-mover effects. Böhme (2005) is concerned that the resulting dominance of one or a few firms may increase vulnerability due to a “monoculture” effect.

¹ In a dynamic context, reputation effects may mitigate some of the externalities – see section “Externalities in a dynamic context.” Furthermore, the recent behavioral economic literature has revealed that economic agents actually take third parties into account, an aspect that will be further studied in the second part of our research project.

Table 1 Externalities as seen from the source of the effect


	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		Volume of malevolent traffic	Risk of proliferating attack	Risk of proliferating attack	Risk of proliferating attack	Hacking opportunities
Large firms		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		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		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

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 substantial spare network capacity will experience only incremental cost increases, again limiting the incentive to invest into ISP-based security upgrades to reduce malware-related traffic.

Security investments by firms, whether they be large corporate users or small and medium-sized 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 equal, will attempt to invest into 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 optimum with precision but they will try to approximate it. In any case, neither the negative external effects of investment falling short of this private optimum nor the positive externalities of investment that goes beyond that optimum are taken into consideration. Individual firm decisions will thus systematically deviate from a social optimum that takes these interdependencies into account.

Figure 1: Externalities with reputation effects



- 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 .

Individual and small business users are by many seen as one of the weakest links in the value chain of networked computing (Camp, unpublished). 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 primarily consider their own costs and benefits but not those of other users. Individual users are thus particularly susceptible to non-intrusive forms of malware, which do not use up significant resources on the user end (e.g., computing power, energy) but in the aggregate 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.

Forms of externalities in networked computer environments

The literature on information security discusses several forms of externalities, including direct 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 causing breaches to spread) 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. They also include the costs associated with the reduced willingness of consumers to engage in e-commerce (Camp and Wolfram 2004).

Externalities in a dynamic framework

In networked computer environments with a rapid pace of technological change, externalities need to be understood in a dynamic framework. With few exceptions, such as Choi et. al. (2005), these aspects are rarely modeled in the present literature. Most importantly, in such a view learning and reputation effects need to be considered. Reputation and learning may happen at different time scales and with different intensity. In any case, they may counteract and reduce the magnitude of negative externalities and possibly enhance positive externalities. Moreover, the incentives and measures of firms to disclose vulnerabilities will influence the magnitude of externalities.

Figure 1 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 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, they externalize costs onto other firms j (C_{ji}). If these cost 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 internalized 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. However, even in this case a reputation effect may work.

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

Empirical estimates of the magnitude of externalities

There are very few empirical estimates of the magnitude of such externalities. For the U.S., a lower boundary is the data collected by the FBI. In 2006, in a sample of 300 reporting firms, the

cost of vulnerabilities was estimated at nearly \$170,000 per year on average (Gordon et al. 2006). Estimates for the total cost to the U.S. economy point to lower two-digit \$ billion figures.

VI. Solutions

A variety of measures have been suggested to internalize externalities. We have discussed some of them in the section on the type of problems that are addressed. Without claiming to be exhaustive, we discuss several solutions that feature prominently in the literature.

Insurance

In recent years, actual insurance policy for cybersecurity risks have come into existence, even though premiums tend to be high because of insufficient data to accurately quantify risk and loss potential (Traub and Leff 2003; Kesan et al. 2005). Premiums can range from \$5,000 to \$60,000 per \$1 million of coverage, depending on the type of business and the extent of insurance coverage (Böhme 2005). Right now, however, the cyber-insurance market is both underdeveloped and underused (Anderson and Moore 2006). Part of this is because worms and viruses are interdependent risks which affect many organizations at the same time, which makes these risks unattractive for insurers. Ogut et al. (2005) are not too optimistic about the effectiveness of insurance when risk is interdependent, even when the market would mature. The same incentives that cause firms to underinvest in security, cause them to buy less insurance coverage. Their model predicts better results from more conventional strategies, such as liabilities and information sharing.

Liability

Liability is an intriguing, but controversial topic. Some authors, like Schneier (2004), see it as the key element of any meaningful strategy for cybersecurity. When applying it to the mushrooming problem of ‘phishing’ (Schneier 2005), he argued: “Push the responsibility – all of it – for identity theft onto the financial institutions, and phishing will go away. This fraud will go away not because people will suddenly get smart and quit responding to phishing e-mails, because California has new criminal penalties for phishing, or because ISPs will recognize and delete the e-mails. It will go away because the information a criminal can get from a phishing attack won’t be enough for him to commit fraud – because the companies won’t stand for all those losses.”

The controversial nature of liability is nowhere more apparent than in the debate over assigning liability to software manufacturers. All proposals are fraught with difficulties. Liability can provide an incentive to produce more secure software, but it can also act as a deterrent to development and innovation. It may also raise entry barriers for new competitors, as large existing companies will find it easier to influence and then to comply with the standard practices that limit their liability.

In the U.S., as elsewhere, there is already legislation that creates liabilities, such as the Gramm-Leach-Bliley (GLB) Act of 1999. GLB includes the ‘Safeguards Rule,’ which requires financial institutions to develop a written information security plan that describes how the company is prepared for, and plans to continue to protect clients’ nonpublic personal information. Notwithstanding such legislation, recent field work suggests that many firms may not see a lack of cybersecurity on their part as a cause for liability issues (Dynes et al. 2005).

Information sharing

Information sharing is an important part of current initiatives for Critical Infrastructure Protection. Several researchers have used models to assess the effects of information sharing on the information security of firms. Gal-Or and Ghose (2005) found that firms have strong incentives to participate in information sharing on security breaches and that this leads to higher security investments, especially in competitive markets and for large organizations. The finding that there are strong incentives to participate is different from the conclusion of Gordon et al. (2003). Their model predicted that information sharing itself would be vulnerable to free riding behavior. Without appropriate incentives, this would prevent information sharing from realizing the potential to reduce overall information security costs and raise social welfare. In more recent work, Gordon and Loeb (2006) conclude that current initiatives of critical infrastructure protection (CIP) lack the appropriate incentive structures.

Markets

A number of authors have proposed to design new markets to internalize current externalities. We already mentioned the work that has been done on a market for vulnerabilities. This market already exists. Currently, two security firms (Tipping Point and iDefense) buy, against undisclosed prices, vulnerabilities and use them to provide security advice for their clients at the same time as they notify the vendor. That way, their clients can update their security measures before the vulnerability becomes public.

Other market designs have also been proposed. Camp and Wolfram (2004) developed the idea of a market for vulnerability credits to be allocated to all machines, analog to emission trading. Ozment (2004) proposes a ‘bug auction,’ which offers a time-variable reward to free-market testers who identify vulnerabilities. A rather unorthodox design was forwarded by McHugh and Deek (2005). They suggest building a ‘microcosm,’ a small-scale, isolated version of the internet. The sponsors of the microcosm would reward any challenger whose malware succeeded in seriously disrupting it. Not everyone is convinced that a vulnerability market would be superior to other solutions. The model of Kannan and Telang (2004) suggests that a market for vulnerabilities almost always underperforms a passive CERT-type mechanism – though the most optimal solution, they argue, would be to let CERT fund vulnerability discovery.

Schechter (2004) advanced the innovative idea that a vulnerability market could also provide us with a reliable metric for the security of a piece of software: If a successful attack is only as difficult as it is to obtain a vulnerability to exploit, then the security strength of that system is

measured by the market price of such a vulnerability. Others have also worked on the issue of how consumers can tell secure from less secure software. This is a classic signaling problem. In the absence of adequate market signals, the information asymmetry between users and manufacturer gives software characteristics of a market for “lemons”, where insecure software may drive out more secure software (Anderson and Moore 2006). Shostack (2005) sees effective signaling as a way to avoid assigning software liabilities, which he views as rather unrealistic and which may also block entry to the market. He discusses a number of possible signals that consumers might use, but finds that all of them are problematic. Auditing tools and certification schemes do not provide adequate signals. In some cases they are actually the opposite, as Edelman (2006) explained. He has shown that certification of trusted sites suffer from adverse selection: The sites that seek and obtain trust certifications are actually significantly less trustworthy than those that forego certification. As the certifying organizations do not bear the consequences of false certification, they have little incentive to invest in more effective verification procedures.

Patching practices is another area in which the design of new markets is explored. When individual users do not apply software patches in a timely fashion, they are vulnerable to attack, but they also increase the risk of attacks for other users – a classic example of an externality. August and Tunca (2006) looked at ways to change the user incentives for patching, finding that in many cases rebates for patching were effective. The software vendor would provide a rebate, where users are compensated by the vendor when a patch is available and they actually patch. This is attractive for the user, but also for the vendor. Their model shows that as more users apply patches, the security of the software increases and thus its value.

Insurance, liability, information sharing and new markets have received a lot of attention in the literature. Of course, there are more ways to internalize externalities: Fines and subsidies, regulation and inspections, technological innovations, voluntary standards or self-regulation, tax policy, government procurement, and governmental standards – to name some of them. While the literature recognizes this, until now there have been no attempts to make systematic comparisons across a wider set of solutions.

VII. Toward a generalized approach

Making cyber space more secure is a technological issue and an issue of influencing people’s behavior. Influencing behavior has to do with many aspects of the human being: psychological, sociological and economical.

Economics is about efficiency, about material and pecuniary incentives to stimulate efficient behavior (in a cost and allocative sense): why do some people buy equipment that is well protected for a higher price than equipment that is less well protected, but is priced lower? Why do firms innovate in software that protects the consumer better? Economics is a social science aiming at understanding causalities from an economic perspective in order to formulate policy recommendations. The science aims to contribute to the realization of specific objectives (e.g., less malware) through designing a specific environment that produces incentives that make actors behave according the preferred outcomes of the processes. The environment in which actors

operate consists of values, norms, laws and regulations (the so-called institutional environment), a technological environment, as well as markets with specific degrees of competition (so-called market structures). All three elements of the institutional environment, the technology and the market structure condition behavior of actors, but at the same time the conditions can be (partly) influenced by them.

Actors in such a specific environment of institutions and market structures act and react according to specific characteristics of the actor: type of rationality (maximization of utility), degree of opportunism, rule following behavior, according to specific norms (shared mental maps).

To understand how actors in a specific market (producers of software, consumers of hardware, criminals in breaking codes) will react to specific technological and social-political measures (laws, regulations, campaigns), demands a clear picture of the 'inner psychology' of the actor, the environment in which he or she is embedded and the latitude for the actors to influence the environment.

Economic theories differ in how actors are modeled (rationality and rule of behavior), what kind environment is taken into account (institutions, technology and market structures) and how the relationship between the actors and the environment is modeled (one directional or interaction).

Mainstream economics (both neoclassical and new institutional economics) models actors with full or bounded rationality and one rule of behavior (maximization of profits and utility, or minimization of costs). In neoclassical economics a firm is a production function positioned in a specific market structure (like pure and perfect competition, or monopolistic competition) forcing the 'actor' to make one optimal combination of production factors given the technology, the institutional environment and the market structure. All actors have the same rather limited inner psychology; they receive the same information through the price system and react in an identical way. Theory is based on axiomatic deductive reasoning aiming at the prediction of optimal end states.

New Institutional economics (NIE) is based on the same type of modeling: methodological individualism – i.e., the actor with homogeneous attributes and one decision rule is constrained by the environment, which is a given to him or her except for the one variable that the theory aims at explaining. In the case of Agency theory that is the optimal contract between the principal and the agent, in case of Transaction Cost Economics that is the optimal mode of governance (institutional arrangements of contracts and organizations that coordinate behavior). The institutional environment as well as the technology are exogenous and given to the actors, the 'institutional arrangements' in the markets are endogenous. New Institutional Economics aims to explain which modes of governance actors will be created. The question to be answered concerns optimal types of contracts and organizations (institutional arrangements) that minimize transaction costs. In NIE there is much attention for the role of property rights (incentives), relationships between principals and agents (producers and customers; shareholder and manager) and transaction costs (optimal contracts and organizations).

In mainstream economics the concept of externalities and the role of government in market economies are well defined. When externalities occur because actions of producers and

consumers cause costs or benefits to others that are not reflected in the prices of the goods and services, then those externalities should first of all be internalized through private contracting. Private actors at micro level know best their preferences and how much money they are willing to pay to others to change the use of their property rights. Externalities should be internalized by voluntary contracting between private actors. Government should create markets if these do not exist (i.e. market for emission rights), and should regulate the markets according to the rules of transparency and enforcements of rights, that the NIE has so well spelled out.

Only in the exceptional case that the negotiation between private actors to internalize an externality is so complex and/ or so costly that the transaction is not performed, government should internalize the externality. This can be done through the well known mechanism of taxes and subsidies, or by prohibiting certain uses of property rights by law.

Most of the economics literature about cybersecurity belongs to the mainstream and especially to the NIE: the insights of advanced microtheory (property rights-, agency- and transaction costs theory, game theory and industrial organization) are frequently applied. This offers useful insights for policy measures aimed at reducing cyber crime. For example the suggestion to change liabilities clearly falls in the category of changing the economic incentives through changing the ‘rules of the game’.

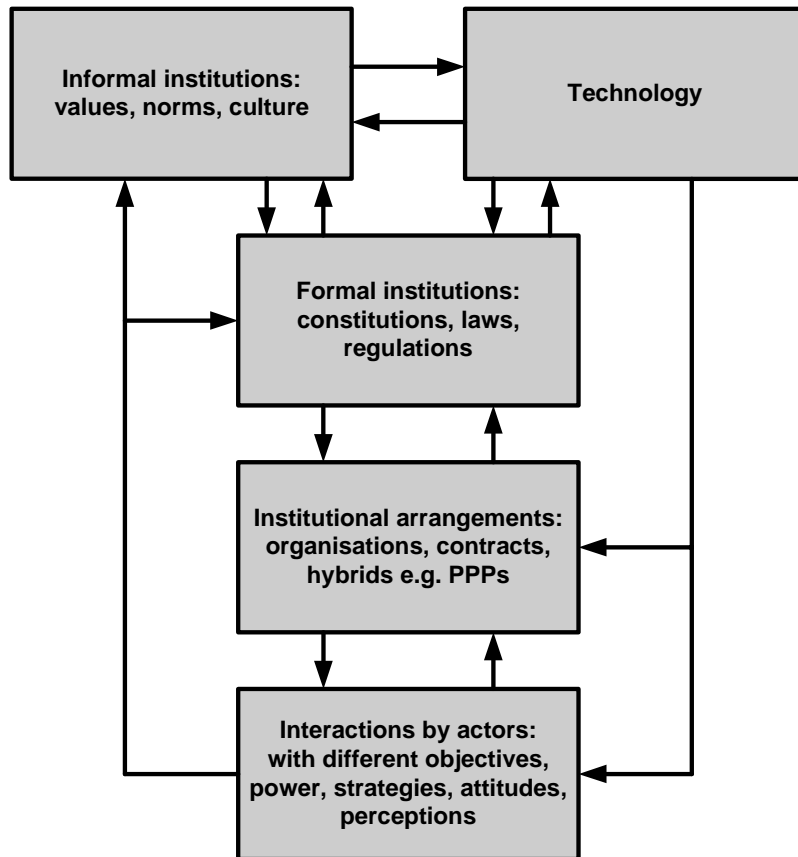
Relevancy of different economic perspectives

Above we have briefly outlined the characteristics of mainstream economics. Theorizing is a matter of isolation: which variables are to be explained, which variables are explanatory and which ones are ignored? A trade-off exists between ‘rigorousness’ of analysis and relevancy. In principle all theorists – and certainly economists – aim at making theories as simple as possible. The point raised here is about the relevancy: are the abstractions about the actors, their environments and the interactions adequate for the issues discussed above? Does the theory correspond sufficiently with the reality of the world of malware and the questions policymakers are interested in?

About processes and dynamics

Mainstream economics is developed to answer optimization questions of a comparative static nature. It basically asks: If the present solution is an equilibrium and one of the parameters (technology, rules of the game, preferences) changes exogenously, what then will be the new equilibrium? About the process of institutionalization mainstream economics assumes that individual actors will adapt or create new institutions, if the existing institutions are not efficient anymore. Government assists with maintaining the rules of the game of competition, which is the driving force towards efficient behavior.

Figure 2: A generalized model of institutional choice and design



It might very well be that the development of adequate policy measures demands an in-depth understanding of the processes of institutional, economic and technological changes that over time interact and influence each other mutually. Then the theoretical framework should be designed to answer questions of dynamics, of processes that emerge, evolve and sometimes break away from existing paths of development to enter a completely new trajectory. In Figure 2 the interrelations are represented: actors of different nature are constraint by institutions, market structures and technology, whereas at the same time these elements of their environment can be (partly) constructed by them. The actors have different objectives, incentive structures and resources to exert power and pressure.

To understand dynamics, co-evolution and punctuated equilibria, frameworks, theories and methods from another paradigm of institutional economics are considered useful. Theories of evolutionary economics, of institutional change related to the so-called Original school of Institutional Economics (OIE) model actors and the interaction with their environment in a different way than the NIE. OIE works with procedural rationality, lock-ins (technological, economical and institutional) and path dependencies, shared mental maps and attempts to capture the dynamics of systems with a framework that is adequate for open systems instead of the closed systems of mainstream economics.

All worlds, including the one of cybersecurity, are dynamic. This issue is: What kind of research questions are we pursuing: comparative static or dynamic questions? A second issue is by what kind of conditions the world of information security is characterized. What kind of rationality do actors have, what kind of markets are the arena for them to act, what about the information (a)symmetry, what about the norms and values ruling in the segment of the world of information security, etc.?

In the world of information security many actors operate. What are their motivations? Cost minimization? Profit maximization? Or follow the rules, norms and habits that prevail in the sector? Do the actors perceive the world around them in the same way, do they operate and react based on the same mental maps? Next to more insights into the motivations of actors, we might need more detailed information about the environment these actors operate in. Different consumers and producers operate in different market structures with different degrees of competition. Moreover, the actors operate in different institutional environments of values and norms, rules and regulations and institutional arrangements (e.g. subject to hacker ethics). The 'inner' and 'outer' environments of the actors determine how they will react to a change in technology or a change in the 'rules of the game'. When an externality emerges what will they do? Will they protect themselves because it is considered their own responsibility? Put pressure on their groups and communities to arrange a collective solution? Will they mobilize politics to change the rules of the game? What will producers do? In case of transparency, the reputation mechanism in the market could be the incentive to improve their products and reduce the externality. What may happen will depend in some cases on a few or, in the extreme case, even on only one variable, but in other cases on a variety of complex interwoven variables. The question how relevant the mainstream of the economics of cyber security is and how much we are in need of insights from OIE depends on the correspondence of the assumptions of the theory with the world we are studying.

VIII. Conclusions

Notwithstanding rapidly growing investments in security measures, it has become clear that cybersecurity is a technological arms race that will not be decided in the immediate future. While many would agree that cybersecurity needs to be strengthened, the effectiveness of many security measures is uncertain and contested, to say the least. Furthermore, security measures may also impede innovation and productivity. Those involved in cybersecurity frequently tend to overlook that the reason why the internet is so susceptible to security threats is the same 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.

All this means that total security is neither achievable nor desirable. Actors make their own tradeoffs regarding what kind of security measures they deem appropriate and rational, given their business model. Clearly, these business models are very different 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 organizations to end users and beyond. In other words, many instances of what could be

conceived as security failures are in fact the outcome of rational economic decisions, given the costs and benefits facing the actors involved.

Our review provides seeks to enhance understanding of the economics of cybersecurity from individual stakeholder perspectives and analyzes their interaction across the value nets of Internet-based commerce. We show that pervasive external effects exist along this value net. Thus, decisions of individual stakeholders are linked in ways that are not reflected in market transactions. Even in the absence of externalities the incentives of individual stakeholders to invest in cybersecurity often deviate from the socially desirable level. The problems are compounded by the organization of the market for cybercrime, which offers many attractive business models for potential criminals, such as low-cost access to botnets. We argue that the dominant neoclassical models of information security, while illustrating important issues, fall short of the complexity of the problem. We offer an alternative framework rooted in a broader institutional approach.

In addition to modeling the economics of malware theoretically, a series of in-depth case studies of key stakeholders in private industry and government was conducted in Europe and the United States. Our findings give a more comprehensive picture than previous attempts, which often focused on a narrow segment of the value chain. We conclude that coping with problems of malware will require measures that reach beyond technological fixes and awareness building. Effective protection will also have to modify the economic incentives of players. Alternative options for restructuring these incentives, from a modification of liability rules to the alteration of the economics of cybercrime are discussed in the paper.

References

- Acquisti, A., A. Friedman and R. Telang (2006). Is There a Cost to Privacy Breaches? An Event Study. Fifth Workshop on the Economics of Information Security 2006, <http://weis2006.econinfosec.org/docs/40.pdf>.
- Acquisti, A. and J. Grossklags (2004). Privacy and Rationality: Preliminary Evidence from Pilot Data. Third Workshop on the Economics of Information Security. Minneapolis, MN.
- 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. (2002). Maybe we spend too much?---Unsettling Parallels Between Security and the Environment. First Annual Workshop on Economics and Information Security. University of California, Berkeley.
- Anderson, R. (2003). Cryptology and Competition Policy-Issues with 'Trusted Computing', http://www.cpppe.umd.edu/rhsmith3/papers/Final_session1_anderson.pdf. 2nd Annual Workshop on Economics and Information Security.
- Anderson, R. and T. Moore (2006). "The Economics of Information Security." Science **314**: 610-613.
- Arora, A., R. Krishnan, A. Nandkumar, R. Telang and Y. Yang (2004). Impact of Vulnerability Disclosure and Patch Availability -- An Empirical Analysis Third Workshop on the Economics of Information Security. Minneapolis, MN.
- August, T. and T. I. Tunca (2006). "Network Software Security and User Incentives." Management Science **52**(11): 1703–1720.
- Bächer, P., T. Holz, M. Kötter and G. Wicherski. (2005). "Know your Enemy: Tracking Botnets. Using honeynets to learn more about Bots." Retrieved 1 March, 2007, from <http://www.honeynet.org/papers/bots/>.
- Bobzin, H. (2006). Principles of Network Economics. Heidelberg, Springer.
- Boehme, R. (2005). Cyber-Insurance Revisited. Fourth Workshop on the Economics of Information Security. Harvard University.
- Böhme, R. (2005). Cyber-Insurance Revisited. Fourth Workshop on the Economics of Information Security. Harvard University.
- Böhme, R. and T. Holz (2006). The Effect of Stock Spam on Financial Markets. Fifth Workshop on the Economics of Information Security 2006, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=897431#PaperDownload.
- Camp, L. J. (2006a). Reliable, Usable Signaling to Defeat Masquerade Attacks. Fifth Workshop on the Economics of Information Security 2006, <http://weis2006.econinfosec.org/docs/48.pdf>.
- Camp, L. J. (2006b). "The State of Economics of Information Security " I/S A Journal of Law and Policy in the Information Society **2**(2): <http://www.is-journal.org/V02I02/2ISJLP189-Camp.pdf>.
- Camp, L. J. (Unpublished). Mental Models of Privacy and Security IEEE Technology and Society.
- Camp, L. J. and S. Lewis, Eds. (2004). Economics of Information Security. Dordrecht, Kluwer Academic Publishers.
- Camp, L. J. and C. Wolfram (2004). Pricing Security: Vulnerability as Externalities: <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.
- Cavusoglu, H., H. Cavusoglu and S. Raghunathan (2005). Emerging issues in responsible vulnerability disclosure Fourth Workshop on the Economics of Information Security. Harvard University.
- Cavusoglu, H., B. Mishra and S. Raghunathan (2004a). "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.
- Cavusoglu, H., S. Raghunathan and B. Mishra (2004b). "A Model for Evaluating IT Security Investments." Communications of the ACM **47**(7): 87-92.
- Chen, P.-y., G. Kataria and R. Krishnan (2005). Software Diversity for Information Security. Fourth Workshop on the Economics of Information Security. Harvard University.
- 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.
- Denning, D. (2000). "Reflections on Cyberweapons Controls." Computer Security Journal **16**(4): 43-53.
- 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>.
- 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.
- Edelman, B. (2006). Adverse Selection in Online 'Trust' Certifications. Fifth Workshop on the Economics of Information Security 2006, <http://weis2006.econinfosec.org/docs/10.pdf>.
- FBI (2006). The Case of the Zombie King: Hacker Sentenced for Hijacking Computers for Profit
- Feigenbaum, J., D. Bergemann, S. Shenker and J. M. Smith (2004). Towards an Economic Analysis of Trusted Systems. Third Workshop on the Economics of Information Security. Minneapolis, MN.
- Frieder, L. and J. Zittrain (2007). Spam Works: Evidence from Stock Touts and Corresponding Market Activity, Social Science Research Network Electronic Paper Collection.
- Friedman, L. S. (2002). The Microeconomics of Public Policy Analysis. Princeton, NJ, Princeton University Press.
- Gal-Or, E. and A. Ghose (2003). The Economic Consequences of Sharing Security Information, http://www.cpppe.umd.edu/rhsmith3/papers/Final_session7_galor.ghose.pdf. 2nd Annual Workshop on Economics and Information Security.
- Gal-Or, E. and A. Ghose (2005). "The Economic Incentives for Sharing Security Information." Information Systems Research **16**(2): 186-208.
- Garcia, A. and B. Horowitz (forthcoming). "The Potential for Underinvestment in Internet Security: Implications for Regulatory Policy." Journal of Regulatory Economics: <http://ssrn.com/abstract=889071>.
- Good, N., J. Grossklags, D. Thaw, A. Perzanowski, D. K. Mulligan and J. Konstan (2006). "User Choices and Regret: Understanding Users Decision Process about Consensually acquired Spyware." I/S A Journal of Law and Policy for the Information Society **2**(2): 283-344.
- Gordon, L. A. and M. P. Loeb (2002). "The Economics of Information Security Investment." ACM Transactions on Information and System Security.

- Gordon, L. A. and M. P. Loeb (2006). Managing Cybersecurity Resources: A Cost-Benefit Analysis. New York, McGraw-Hill.
- Gordon, L. A., M. P. Loeb and W. Lucyshyn (2003). "Sharing Information on Computer Systems: An Economic Analysis." Journal of Accounting and Public Policy **22**(6): 461-485.
- Gordon, L. A., M. P. Loeb, W. Lucyshyn and R. Richardson (2006). 2006 CSI/FBI Computer Crime and Security Survey. San Francisco, Computer Security Institute.
- Jan, R. J. (2007). Sophos: US wins spam and malware crown. SC Magazine.
- 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.
- Justice, D. o. (2004). Background On Operation Web Snare: Examples Of Prosecutions, Department of Justice.
- Kannan, K. and R. Telang (2004). An economic analysis of market for software vulnerabilities. Third Workshop on the Economics of Information Security. Minneapolis, MN.
- Keizer, G. (2005). From Russia With Malware. InformationWeek
- Kesan, J. P., R. P. Majuca and W. J. Yurcik (2005). Cyber-insurance As A Market-Based Solution To The Problem Of Cybersecurity. Fourth Workshop on the Economics of Information Security. Harvard University.
- Kirk, J. (2007). Microsoft falls victim to shady 'scareware'. InfoWorld.
- Kunreuther, H. and G. Heal (2003). "Interdependent security." Journal of Risk and Uncertainty **26**(2): 231.
- LaRose, R., N. Rifon, S. Liu and D. Lee (2005). Understanding Online Safety Behavior: A Multivariate Model. International Communication Association New York.
- Leyden, J. (2006, 5 December 2006). "Malware wars: Are hackers on top?" The Register, from http://www.theregister.co.uk/2006/12/05/malware_trends/.
- Leyden, J. (2007, 2 February 2007). "Dutch botnet herder fined €75K for sending 9bn spams." The Register, from http://www.theregister.co.uk/2007/02/02/dutch_spammer_fined/.
- Lookabaugh, T. and D. C. Sicker (2003). Security and Lock-In: The Case of the U.S. Cable Industry, http://www.cpppe.umd.edu/rhsmith3/papers/Final_session8_lookabaugh.sicker.pdf. 2nd Annual Workshop on Economics and Information Security.
- MailChannels. (2007). "Spamonomics 2.0: Interrupting the Economics of Spamming with Traffic Shaping." from http://mailchannels.com/documents/spamonomics_whitepaper_Jan10_2007.pdf.
- McHugh, J. A. M. and F. P. Deek (2005). "An incentive system for reducing malware attacks." Association for Computing Machinery. Communications of the ACM **48**(6): 94.
- Mell, P., K. Kent and J. Nusbaum (2005). Guide to Malware Incident Prevention and Handling. Gaithersburg, MD, National Institute of Standards and Technology.
- Nicolai, J. (2000). Analyst Puts Hacker Damage at \$1.2B. InfoWorld.
- Nizovtsev, D. and M. Thursby (2005). Economic Analysis of Incentives to Disclose Software Vulnerabilities. Fourth Workshop on the Economics of Information Security. Harvard University.
- Odlyzko, A. (2004). Privacy, Economics, and Price Discrimination on the Internet. Economics of Information Security. L. J. Camp and C. Wolfram. Dordrecht, Kluwer Academic Publishers: 187-214.

- OECD. (2002a). "Guidelines for the Security of Information Systems and Networks." from <http://www.oecd.org/dataoecd/16/22/15582260.pdf>.
- OECD (2002b). The OECD 2002 Security Guidelines - Q&A.
- Ogut, H., N. Menon and S. Raghunathan (2005). Cyber insurance and IT security investment: Impact of interdependent risk. Fourth Workshop on the Economics of Information Security. Harvard University.
- Ozment, A. (2004). Bug auctions: Vulnerability markets reconsidered. Third Workshop on the Economics of Information Security. Minneapolis, MN.
- Ozment, A. and S. E. Schechter (2006). Milk or Wine: Does Software Security Improve with Age? The Fifteenth Usenix Security Symposium. Vancouver, BC, Canada.
- Png, I. P. L., C. Q. Tang and Q.-H. Wang (2006). Hackers, Users, Information Security. Fifth Workshop on the Economics of Information Security 2006, <http://weis2006.econinfosec.org/docs/54.pdf>.
- 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.
- Poulsen, K. (2003a). Nachi worm infected Diebold ATMs. SecurityFocus.
- Poulsen, K. (2003b). Slammer worm crashed Ohio nuke plant network. SecurityFocus.
- Rescorla, E. (2004). Is finding security holes a good idea? . Third Workshop on the Economics of Information Security. Minneapolis, MN.
- Rifon, N., E. T. Quilliam and R. LaRose (2005). Consumer Perceptions of Online Safety. International Communication Association New York.
- 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, <http://weis2006.econinfosec.org/docs/18.pdf>.
- Schechter, S. E. (2004). Computer Security Strength & Risk: A Quantitative Approach.
- Schneier, B. (2004). Secrets and Lies: Digital Security in a Networked Society. New York, Wiley.
- Schneier, B. (2005). A Real Remedy for Phishers, 6 October 2005. Available online at: <http://www.wired.com/news/politics/0,1283,69076,00.html>. Wired News.
- Shostack, A. (2005). Avoiding Liability: An Alternative Route to More Secure Products. Fourth Workshop on the Economics of Information Security. Harvard University.
- Soo Hoo, K., A. W. Sudbury and A. R. Jaquith (2001). "Tangible ROI through Secure Software Engineering." Secure Business Quarterly 1(2): 1-3.
- Sturgeon, W. (2005). Denial of service attack victim speaks out. SME Director.
- Thompson, C. (2004). The Virus Underground. The New York Times Magazine. New York.
- Tiwari, R. K. and K. Karlapalem (2005). Cost Tradeoffs For Information Security Assurance. Fourth Workshop on the Economics of Information Security. Harvard University.
- Traub, R. K. and R. M. Leff (2003). "Insurance coverage for cyber losses[dagger]." FDCC Quarterly 53(4): 357.
- Trendmicro (2007). The Phishing Ecosystem: The Players And Their Interactions. First Line of Defense Newsletter.
- Varian, H. (2000). "Managing Online Security Risks." New York Times.
- Varian, H. (2004). System Reliability and Free Riding. Economics of Information Security. L. J. Camp and S. Lewis. Dordrecht, Kluwer Academic Publishers: 1-16.
- Vijayan, J. (2003). "Improved security through IT diversity." Computerworld 37(47): 28.

Wathieu, L. and A. Friedman (2005). An Empirical Approach to Understanding Privacy Valuation. Fourth Workshop on the Economics of Information Security. Harvard University.

Weaver, N. and V. Paxson (2004). A Worst-Case Worm. Third Workshop on the Economics of Information Security. Minneapolis, MN.

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		Volume of malevolent traffic	Risk of proliferating attack	Risk of proliferating attack	Risk of proliferating attack	Hacking opportunities
Large firms		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		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		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.