

An Introduction to Security Economics

Richard Clayton
`richard.clayton@cl.cam.ac.uk`

with acknowledgements to

Ross Anderson
`ross.anderson@cl.cam.ac.uk`

&

Tyler Moore
`tylerm@smu.edu`



UNIVERSITY OF
CAMBRIDGE
Computer Laboratory

Part II: Security
30th January 2013

NPL 
National Physical Laboratory

Outline

- Security economics
 - a powerful new way of looking at overall system security
- Some of the key basic ideas from economics
 - incentives
 - asymmetric information
 - externalities
 - adverse selection
- Security economics research examples
 - adverse selection in security seals
 - markets for vulnerabilities
 - phishing website takedown
 - the cost of cybercrime

Traditional View of Information Security

- People used to think that the reason that the Internet was insecure because of a lack of features, or that there was not enough crypto / authentication / filtering
- Plus, 'if only' people had a really good checklist of security issues to tackle, then we would all be more secure
- So engineers worked on providing better, cheaper, (and even occasionally easy-to-use) security features – developing secure building blocks such as SHA-1, AES, PKI, firewalls...
- Others worked on long lists of things to check up on, or policies that ought to be adopted...
- About 1999, we started to realize that this is not enough...

The 'New School' of Information Security

- For the last decade, we have started to apply an economic analysis to information security issues
- Economic analysis often addresses the underlying causes of security failures within a system, whereas a technical analysis will merely identify the mechanism!
- Tackling the problem in economic terms can lead to valuable insights as to how to create permanent fixes
- Clearly shows that consumers need access to better information so they can make informed decisions about security
- Meanwhile, the trend is for information security mechanisms (such as cryptographic protocols) to be used to support business models rather than to manage risk

New Uses of Security Mechanisms

- Xerox started using authentication in ink cartridges to tie them to the printer
 - followed by HP, Lexmark. . . and Lexmark's case against SCC
 - note that the profit is in the consumables – purchasers compare ticket price, rather than total cost of ownership
- Accessory control now spreading to more and more industries
 - games, mobile phones, cars...
- Digital rights management (TPMs): Apple grabs control of music downloads, games consoles almost given away and money is made from licensing deals to allow games to be played...
- Cryptography is being used to tackle the obvious contradiction between the decentralization of network intelligence and the operators desire to retain control

Using Economics to Explain Security

- Electronic banking: UK banks were less liable for fraud than US banks, so they got careless and ended up suffering more fraud and error. The economists call this a 'moral hazard'
- Distributed denial of service: viruses no longer attack the infected machine but they use it to attack others. Why should customers spend \$50 on anti-virus software when it isn't their data that is trashed? Economists call this an 'externality'
- Health records: hospitals, not patients, buy IT systems, so they protect the hospitals' interests rather than patient privacy. These are 'incentive' and 'liability' failures

and

- Why is Microsoft software so insecure, despite its market dominance? The economists can explain this as well!

Security Economics Research

- Key early work by Anderson, Odzlyko & Schneier
- Security Economics has grown to 100+ active researchers
- Workshop on the Economics of Information Security (WEIS), held annually in major research centers in US and UK
- Topics range from econometrics of online crime through DRM policy to return on security investment and how to manage the patching cycle
- Anderson maintains an 'Economics and Security Resource Page'

<http://www.cl.cam.ac.uk/~rja14/econsec.html>

- Note also various survey papers by Anderson & Moore, the latest of which is:

<ftp://ftp.deas.harvard.edu/techreports/tr-03-11.pdf>

The Basics of the New Analysis

- Incentives: failures are more likely when the person responsible for protecting a system is not the one who suffers harm
 - so it's of concern if a bank can dump 'phishing' losses onto customers; or if hospital systems put administrator convenience before patient privacy
- Asymmetric information
 - vendors claim that their software is secure, but the buyers have no means of judging this; so they refuse to pay a premium for quality
- Externalities ('side effects')
 - a larger network is more valuable to each of its members, so there is a trend towards dominance (Microsoft/Facebook/iTunes)
 - 'negative externalities' arise where the damage is done to someone else; malware may not do much local damage, but botnet membership means that everyone else is being damaged

IT Economics and Security I

- The high fixed and low marginal costs, the network effects and switching costs are all powerful drivers towards dominant-firm markets with a big 'first-mover' advantage
- Hence the 'time-to-market' is critical
- Paying attention to security rarely assists scheduling
- Hence the Microsoft philosophy of "we'll ship it Tuesday and get it right by version 3" is not perverse behaviour by Bill Gates, or a moral failing, but absolutely rational behaviour
- If Microsoft had not acted this way, then another company which took this approach would have become the dominant player in the PC operating system business (and/or in the office productivity tools business)

IT Economics and Security II

- When building a network monopoly, it is critical to appeal to the vendors of complementary products
 - remember the old mantra of “find the software product then ask which machine and operating system to buy”...
 - ... Microsoft spent huge amounts assisting developers
 - we can see the same pattern with PC v Apple; Symbian v WinCE, WMP v RealPlayer, not to mention the console games market
- The lack of security in earlier versions of Windows made it significantly easier to develop applications
- It's also easy for vendors to choose security technologies that dump support costs onto the users (SSL not SET, PKI, . . .)
- SSH succeeded because the switching cost was low (Telnet+ +) and there's benefit to early adopters; hence BGPSEC, DNSSEC and various email protection schemes struggle

The Economics 'Rules' for the IT Industry

- Network effects
 - value of a network grows super-linearly to its size (Metcalfe's Law says n^2 , Briscoe/Odlyzko/Tilly suggest $n \log n$)
 - this drives monopolies, and is why we have just one Internet
- High fixed and low marginal costs
 - competition drives price down to marginal costs of production; but in IT industries this is usually (near as makes no difference) zero
 - hence copyright, patents etc. needed to recover capital investment
- Switching costs determine value
 - switching from an IT product or service is usually expensive
 - once you have 1000 songs on your iPod, you're locked into iPods
 - Shapiro-Varian theorem: net present value of a software company is the total switching costs of its customers

Privacy – a puzzling area

- Most people say they value privacy, but act otherwise. Most privacy ventures have failed
- So why is there this privacy gap?
- Hirshleifer – privacy is a means of social organization, a legacy of territoriality
- Odlyzko – technology makes price discrimination both easier and more attractive; and to efficiently discriminate you need to know more about your customers
- Acquisti – people care about privacy when buying clothes, but not cameras (phone viruses worse for image than PC viruses?)
- Leads to research in behavioural economics (the interface between economics and psychology)

Key Problem of the Information Society

- More and more goods contain software so more and more industries are starting to become like the software industry
- The Good
 - flexibility, rapid response
- The Bad
 - complexity, frustration, bugs
- The Ugly
 - attacks, frauds, monopolies
- When markets fail, one way of dealing with this is to regulate, so how will regulation evolve to cope with this?

Adverse Selection & Moral Hazard

- Suppose you sell insurance to smokers and non-smokers. Smokers are more likely to die earlier, so they get better value from insurance than non-smokers, so as a group they buy more insurance – so the insured are a worse risk. From the point of view of the insurance company the higher mortality by those who 'select' insurance is 'adverse'.
 - fix is to require medicals, or use questionnaires to set rates
- Some central bankers did not want to bail out the failing banks because of the 'moral hazard' (the removal of the incentive to be prudent in future)
- Why do Volvo drivers have more accidents? Adverse selection can lead to bad drivers choosing Volvos and moral hazard may mean that people drive more badly because they feel safe
 - (and the "risk thermostat" may mean that drive more recklessly!)

Adverse Selection in Security Software

- George Akerlof's 'market for lemons' (Nobel Prize 2001)
 - considered the trade in second-hand cars as a metaphor for a market with asymmetric information: if there are 50 cars worth \$2K and 50 cars worth \$1K, then what is the equilibrium price?
 - buyers cannot determine car quality, so they are unwilling to pay a premium for a quality car
 - sellers know this, so market is dominated by low-quality goods
- Software market is a market for lemons (Anderson 2001)
 - vendors may believe their software is secure, but buyers have no reason to accept that this is correct
 - so buyers refuse to pay a premium for secure software, and vendors refuse to devote resources to make it secure
- How can we reduce this asymmetry of information?

Adverse Selection in Seals and Adverts

- Ben Edelman (WEIS 2006) used data from SiteAdvisor to identify 'bad' sites distributing spam and malware
 - 2.5% of all sites were found to be 'bad'
- But 'bad' companies are more likely to be TRUSTe-certified:
 - 5.4% of TRUSTe-certified sites are 'bad'
 - However, sites with the BBBOnline seal are slightly more trustworthy than random sites (but their process is very slow and there were only 631 certificates issued)
- Similarly, untrustworthy sites are over-represented in paid advertisement links compared to the organic search results
 - 2 to 3% of organic results are 'bad' (0% for top hit at Yahoo!)
 - 5 to 8% of advertising links are 'bad'

Tackling Adverse Selection by Regulation

- When the market fails you regulate!
- Options:
 - require certification authorities and search engines to devote more resources to policing content
 - assign liability to certification entities if certifications are granted without proper vetting
 - alternatively, regulate enforcement actions by requiring complaints to be published
 - search engine operators could be required to exercise 'reasonable diligence' before agreeing to accept an advertisement
- But so far, we're just tolerating/ignoring the problem

Markets for Vulnerabilities

- Need a way to easily measure a system's security
 - stock markets dip after breach, but only a bit & soon forgotten
- One possible approach: establish a market price for an undiscovered vulnerability (Schechter 2002)
 - reward software testers (hackers) for identifying new vulnerability
 - products with higher outstanding rewards are more secure
- Not simply academic fantasy
 - iDefense, Tipping Point have created quasi-markets for vulnerabilities (& WabiSabiLabi had an auction site for a while)
 - however, their business models have been shown to be socially sub-optimal (e.g., they provide disclosure information only to subscribers and they have an incentive to disclose vulnerabilities to harm non-subscribers)
 - limited public information (at present) on pricing

Malware on the Internet

- Internet security suffers from negative externalities
- Modern malware harms others far more than its host: botnet machines send spam and do all the other bad things, but the malware doesn't usually trash the disk and may try to avoid over-use of bandwidth or processing cycles
- ISPs find quarantine and clean-up expensive (an interaction between customer and helpdesk costs more than the profit from that customer for months to come)
- ISPs are not harmed much by insecure customers since it's just a bit more traffic and a handful of complaints to process
- Should the Government have a role here (c.f. the way in which we tackle illness by public health initiatives)
 - the debate on this tells you more about participants' political views than whether this is a valuable suggestion

Takedown Times: Moore/Clayton WEIS 08

- Defamation – believed to be quick (days)
 - copyright violation – also prompt(ish)
 - experimentally 'days' (with prompting, so perseverance matters)
- Fake escrow agents
 - average 9 days, median 1 day
- Phishing
 - 4 hours if bank aware, 4 days if not
- Mule recruitment sites (Sydney Car Center etc.)
 - average 13 days, median 8 days
 - doesn't attack any particular bank, so they ignore the issue
 - slower than escrow sites (vigilantes more motivated ?)
- Fake pharmacies
 - no 'vigilante groups' – so lifetime is ~2 months

Measuring Cybercrime

- 2009 McAfee: cybercrime costs \$1000bn (1 trillion) worldwide
- 2011 Detica (part of BAE plc): estimated cost of cybercrime to the UK economy was \$43 billion / annum (~ 1.8% of GDP)
- Florencio and Herley "Sex, Lies and Cybercrime Surveys"
 - this WEIS 2011 paper points out how outliers affect results (single loss of \$50K in a 1000 person survey becomes \$10bn scaled up)
- We (multiple expert authors) assessed data for WEIS 2012:
 - created framework, and gave best estimates for each category
 - traditional frauds cost citizens a few hundred dollars per year
 - transitional frauds cost citizens a few tens of dollars per year
 - new cybercrimes net criminals tens of pence per citizen per year
- BUT the indirect costs and defence costs (and especially cleanup costs) for new crimes are more than 10x the criminal revenue

The Research Agenda

- The online world and the physical world are merging, and this will cause major dislocation for many years to come
- Security economics gives us some of the tools we need to understand what's going on
 - we're a lot less puzzled than we were in 2000!
- Sociology gives some cool and useful stuff too
- And a new focus on 'security psychology' is not just about usability and preventing phishing. It might bring us fundamental insights, particularly in improving our understanding of why security fails for some individuals – just as security economics has given us insight into why it can fail for the crowd

More..

Economics and Security Resource Page

<http://www.cl.cam.ac.uk/~rja14/econsec.html>

Cambridge Security Group Blog

<http://www.lightbluetouchpaper.org>

