Information Security Economics – and Beyond

Ross Anderson and Tyler Moore

Computer Laboratory, University of Cambridge
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

Abstract

The economics of information security has recently become a thriving and fast-moving discipline. As distributed systems are assembled from machines belonging to principals with divergent interests, incentives are becoming as important to dependability as technical design. The new field provides valuable insights not just into 'security' topics such as privacy, bugs, spam, and phishing, but into more general areas such as system dependability (the design of peer-to-peer systems and the optimal balance of effort by programmers and testers), and policy (particularly digital rights management). This research program has been starting to spill over into more general security questions (such as law-enforcement strategy), and into the interface between security and the social sciences. Most recently it has started to interact with psychology, both through the psychology-and-economics tradition and in response to phishing. The promise of this research program is a novel framework for analyzing information security problems – one that is both principled and effective.

1 Introduction

Since about 2000, people have started to realise that security failure is caused by bad incentives at least as often as by bad design. Systems are particularly prone to failure when the person guarding them does not suffer the full cost of failure. Game theory and microeconomic theory are becoming important to the security engineer, just as as the mathematics of cryptography did a quarter century ago. The growing use of security mechanisms for purposes such as digital rights management and accessory control – which exert power over system owners rather than protecting them from outside enemies – introduces many strategic issues. Where the system owner's interests conflict with those of her machine's designer, economic analysis can shine light on policy options.

We survey recent results and live research challenges in the economics of information security. Our goal is to present several promising applications of economic theory and ideas to practical information security problems. In Section 2, we consider foundational concepts: misaligned incentives in the design and deployment of computer systems, and the impact of externalities. Section 3 discusses information security applications where economic analysis has yielded interesting insights: software vulnerabilities, privacy, and the development of user-control mechanisms to support new business models. Metrics present another challenge: risks cannot be managed better until they can be measured better. Most users cannot tell good security from bad, so developers are not compensated for efforts to strengthen their code. Some evaluation schemes are so badly managed that 'approved' products are less secure than random ones. Insurance is also problematic; the

local and global correlations exhibited by different attack types largely determine what sort of insurance markets are feasible. Cyber-risk markets are thus generally uncompetitive, underdeveloped or specialised.

Economic factors also explain many challenges to privacy. Price discrimination – which is economically efficient but socially controversial – is simultaneously made more attractive to merchants, and easier to implement, by technological advance. Privacy problems also create many externalities. For example, spam and 'identity theft' impose non-negligible social costs. Information security mechanisms or failures can also create, destroy or distort other markets: digital rights management in online music and software markets provides a topical example. Finally, we look at government policy options for dealing with market failures in Section 4, where we examine regulation and mechanism design.

We conclude by discussing several open research challenges: examining the security impact of network structure on interactions, reliability and robustness.

2 Foundational Concepts

Economic thinkers used to be keenly aware of the interaction between economics and security; wealthy nations could afford large armies and navies. But nowadays a web search on 'economics' and 'security' turns up relatively few articles. The main reason is that, after 1945, economists drifted apart from people working on strategic studies; nuclear weapons were thought to decouple national survival from economic power [1], and a secondary factor may have been that the USA confronted the USSR over security, but Japan and the EU over trade. It has been left to the information security world to re-establish the connection.

2.1 Misaligned incentives

One of the observations that sparked interest in information security economics came from banking. In the USA, banks are generally liable for the costs of card fraud; when a customer disputes a transaction, the bank must either show she is trying to cheat it, or refund her money. In the UK, the banks had a much easier ride: they generally got away with claiming that their systems were 'secure', and telling customers who complained that they must be mistaken or lying. "Lucky bankers," one might think; yet UK banks spent more on security and suffered more fraud. This may have been what economists call a moral-hazard effect: UK bank staff knew that customer complaints would not be taken seriously, so they became lazy and careless, leading to an epidemic of fraud [2].

In 1997, Ayres and Levitt analysed the Lojack car-theft prevention system and found that once a threshold of car owners in a city had installed it, auto theft plummeted, as the stolen car trade became too hazardous [3]. This is a classic example of an *externality*, a side-effect of an economic transaction that may have positive or negative effects on third parties. Camp and Wolfram built on this in 2000 to analyze information security vulnerabilities as negative externalities, like air pollution: someone who connects an insecure PC to the Internet does not face the full economic costs of that, any more than someone burning a coal fire. They proposed trading vulnerability credits in the same way as carbon credits [4].

Also in 2000, Varian looked at the anti-virus software market. People did not spend as much on protecting their computers as they logically should have. At that time, a typical virus payload was a service-denial attack against the website of Amazon or Microsoft. While a rational consumer might well spend $20 to stop a virus trashing her hard disk, she will be less likely to do so just to protect a wealthy corporation [5].

Legal theorists have long known that liability should be assigned to the party that can best manage the risk. Yet everywhere we look, we see online risks allocated poorly, resulting in privacy failures and protracted regulatory tussles. For instance, medical record systems are bought by hospital directors and insurance companies, whose interests in account management, cost control and research are not well aligned with the patients' interests in privacy. This mismatch of incentives led in the USA to HIPAA, a law that sets standards for privacy in health IT. Bohm et al. [6] documented how many banks used online banking as a means of dumping on their customers many of the transaction risks that they previously bore in the days of cheque-based banking; for an update on liability in payment systems, see [7].

Asymmetric information plays a large role in information security. Moore showed that we can classify many problems as hidden-information or hidden-action problems [8]. The classic case of hidden information is the 'market for lemons' [26]. Akerlof won a Nobel prize for the following simple yet profound insight: suppose that there are 100 used cars for sale in a town: 50 well-maintained cars worth $2000 each, and 50 'lemons' worth $1000. The sellers know which is which, but the buyers don't. What is the market price of a used car? You might think $1500; but at that price no good cars will be offered for sale. So the market price will be close to $1000. Hidden information, about product quality, is one reason poor security products drive out good ones. When users can't tell good from bad, they might as well buy a cheap antivirus product for $10 as a better one for $20, and we may expect a race to the bottom on price.

Hidden-action problems arise when two parties wish to transact, but one party's unobservable actions can impact the outcome. The classic example is insurance, where a policyholder may behave recklessly without the insurance company observing this. We find the same general problem in networks. Network nodes can hide malicious or antisocial behavior from their peers; routers can quietly drop selected packets or falsify responses to routing requests; nodes can redirect network traffic to eavesdrop on conversations; and players in file-sharing systems can hide whether they share with others, so some may 'free-ride' rather than to help sustain the system. Once the problem is seen in this light, designers can minimise the capacity for hidden action, or to make it easy to enforce suitable contracts.

This helps explain the evolution of peer-to-peer systems. Early systems proposed by academics, such as Eternity, Freenet, Chord, Pastry and OceanStore, required users to serve a random selection of other users' files [9]. These systems were never widely adopted. Later systems that did attract large numbers of users, like Gnutella and Kazaa, instead allow peer nodes to serve only the content they have downloaded for their own use, rather than burdening them with others' files. The comparison between these architectures originally focused on purely technical aspects: the cost of search, retrieval, communications and storage. However, analysing incentives turned out to be fruitful too.

First, a system structured as an association of clubs reduces the potential for hidden action; club members are more able to assess which members are contributing. Second, clubs might have quite divergent interests. Though peer-to-peer systems are now seen as mechanisms for sharing music, early systems were designed for censorship resistance. A system might serve a number of quite different groups – maybe Chinese dissidents, critics of Scientology, or aficionados of sado-masochistic imagery that is legal in California but banned in Tennessee. Early peer-to-peer systems required such users to serve each other's files, so that they ended up protecting each others' free speech. But might such groups not fight harder to defend their own colleagues, rather than people involved in struggles in which they have no interest?

Danezis and Anderson introduced the Red-Blue model to analyze this [10]. Each node has a preference among resource types, for instance left-leaning versus right-leaning political texts, while a censor will try to impose his own preference. His action will suit some nodes but not others. The model proceeds as a multi-round game in which nodes set defense budgets that affect the probability that they will defeat the censor or be overwhelmed by him. Under reasonable assumptions, the authors show that diversity (with each node storing its preferred resource mix) performs better under attack than solidarity (where each node stores the same resource mix). Diversity makes nodes willing to allocate higher defense budgets; the greater the diversity, the more quickly will solidarity crumble in the face of attack. This model was an early venture on the boundary between economics and sociology; it sheds light on the general problem of diversity versus solidarity, which has had a high profile recently because of the question whether the growing diversity of modern societies is in tension with the solidarity on which modern welfare systems are founded [11].

2.2 Security as an externality

Information industries have many different types of externality. They tend to have dominant firms for three reasons. First, there are often network externalities, whereby the value of a network grows more than linearly in the number of users; for example, anyone wanting to auction some goods will usually go to the largest auction house, as it will attract more bidders. Second, there is often technical lock-in stemming from interoperability, and markets can be two-sided: software firms develop for Windows to access more customers, and users buy Windows machines to get access to more software. Third, information industries tend to combine high fixed and low marginal costs: the first copy of a software program (or a DVD) may cost millions to produce, while subsequent copies are almost free. These three features separately can lead to industries with dominant firms; together, they are even more likely to.

This not only helps explain the rise and dominance of operating systems, from System/360 through DOS and Windows to Symbian; it also helps explain patterns of security flaws. While a platform vendor is building market dominance, it has to appeal to vendors of software as well as to users, and security could get in their way. So vendors start off with minimal protection; once they have become dominant, they add security to lock their customers in more tightly [12]. We'll discuss this in more detail later.

Further externalities affect security investment, as protection often depends on the efforts of many principals. Hirshleifer told the story of Anarchia, an island whose flood defences

were constructed by individual families and whose defence depends on the weakest link, that is, the laziest family; he compared this with a city whose defences against ICBM attack depend on the single best defensive shot [13]. Varian extended this to three cases of interest to the dependability of information systems – where performance depends on the minimum effort, the best effort, or the sum-of-efforts [14].

Program correctness can depend on minimum effort (the most careless programmer introducing a vulnerability) while software vulnerability testing may depend on the sum of everyone's efforts. Security may also depend on the best effort – the actions taken by an individual champion such as a security architect. When it depends on the sum of individual efforts, the burden will tend to be shouldered by the agents with the highest benefit-cost ratio, while the others free-ride. In the minimum-effort case, the agent with the lowest benefit-cost ratio dominates. As more agents are added, systems become more reliable in the total-effort case but less reliable in the weakest-link case. What are the implications? Well, software companies should hire more software testers and fewer but more competent programmers. (Of course, measuring programmer competence can be hard, which brings us back to hidden information.)

This work inspired other researchers to consider interdependent risk. A recent influential model by Kunreuther and Heal notes that an individual taking protective measures creates positive externalities for others that in turn may discourage them from investment [15]. This insight has implications far beyond information security. The decision by one apartment owner to install a sprinkler system will decrease his neighbours' fire risk and make them less likely to do the same; airlines may decide not to screen luggage transferred from other carriers who are believed to be careful with security; and people thinking of vaccinating their children may choose to free-ride off the herd immunity instead. In each case, several widely varying equilibria are possible, from complete adoption to total refusal, depending on the levels of coordination between principals.

Katz and Shapiro famously analyzed how network externalities influenced the adoption of technology: they lead to the classic S-shaped adoption curve in which slow early adoption gives way to rapid deployment once the number of users reaches some critical mass [16]. Network effects can also influence the initial deployment of security technology, whose benefit may depend on the number of users who adopt it. The cost may exceed the benefit until a minimum number adopt; so everyone might wait for others to go first, and the technology never gets deployed. Recently, Ozment and Schechter have analyzed different approaches for overcoming such bootstrapping problems [17].

This challenge is particularly topical. A number of core Internet protocols, such as DNS and routing, are considered insecure. Better protocols exist (e.g., DNSSEC, S-BGP); the challenge is to get them adopted. Two widely-deployed security protocols, SSH and IPsec, both overcame the bootstrapping problem by providing significant internal benefits to adopting firms, with the result that they could be adopted one firm at a time, rather than needing everyone to move at once. The deployment of fax machines was similar: many companies initially bought fax machines to connect their own offices.

3 Applications

3.1 Economics of vulnerabilities

There has been much debate about 'open source security', and more generally whether actively seeking and disclosing vulnerabilities is socially desirable. Anderson showed in 2002 that, under standard assumptions of reliability growth, open systems and proprietary systems are just as secure as each other; opening up a system helps the attackers and defenders equally [18]. Thus the open-security question may be an empirical one, turning on the extent to which a given real system follows the standard model.

Rescorla argued in 2004 that for software with many latent vulnerabilities, removing one bug makes little difference to the likelihood of an attacker finding another one later [19]. Since exploits are often based on vulnerabilities inferred from patches, he argued against disclosure and frequent patching unless the same vulnerabilities are likely to be rediscovered. This raised the question of whether software follows the standard dependability model, of independent vulnerabilities. Ozment found that for FreeBSD, vulnerabilities are correlated in that they are likely to be rediscovered [20]. Ozment and Schechter also found that the rate at which unique vulnerabilities were disclosed for the core and unchanged FreeBSD operating system has decreased over a six-year period [21]. These findings suggest that vulnerability disclosure can improve system security over the long term. Vulnerability disclosure also helps motivate vendors to fix bugs [22]. Arora et al. showed that public disclosure made vendors respond with fixes more quickly; attacks increased to begin with, but reported vulnerabilities declined over time [23].

This discussion begs a deeper question: why do so many vulnerabilities exist in the first place? A useful analogy might come from considering large software project failures: it has been known for years that perhaps 30% of large development projects fail [24], and this figure does not seem to change despite improvements in tools and training: people just built much bigger disasters nowadays than they did in the 1970s. This suggests that project failure is not fundamentally about technical risk but about the surrounding socio-economic factors (a point to which we will return later). Similarly, when considering security, software writers have better tools and training than ten years ago, and are capable of creating more secure software, yet the economics of the software industry provide them with little incentive to do so.

In many markets, the attitude of 'ship it Tuesday and get it right by version 3' is perfectly rational behaviour. Many software markets have dominant firms thanks to the combination of high fixed and low marginal costs, network externalities and client lock-in noted above [25], so winning market races is all-important. In such races, competitors must appeal to complementers, such as application developers, for whom security gets in the way; and security tends to be a lemons market anyway. So platform vendors start off with too little security, and such as they provide tends to be designed so that the compliance costs are dumped on the end users [12]. Once a dominant position has been established, the vendor may add more security than is needed, but engineered so as to maximise customer lock-in [27].

In some cases, security is even worse than a lemons market: even the vendor does not know how secure its software is. So buyers have no reason to pay more for protection, and vendors are disinclined to invest in it.

How can this be tackled? Economics has suggested two novel approaches to software security metrics: vulnerability markets and insurance.

Vulnerability markets help buyers and sellers establish the actual cost of finding a vulnerability in software. To begin with, some standards specified a minimum cost of various kinds of technical compromise; one example is banking standards for point-of-sale terminals [28]. Another example comes from bug bounties: Don Knuth offered rewards to people who found bugs in his typesetting software while the Mozilla Foundation offers $500 for critical security vulnerabilities in Firefox. Camp and Wolfram suggested in 2000 that markets might work better here [4], and Schechter developed this into a proposal for open markets in reports of previously undiscovered vulnerabilities [29]. Two firms, iDefense and Tipping Point, are now openly buying vulnerabilities, so the market actually exists (unfortunately, the prices are not published). Their business model is to provide vulnerability data simultaneously to their customers and to the affected vendor, so that their customers can update their firewalls before anyone else. However, the incentives here are suboptimal: bug-market organisations might increase the value of their product by leaking vulnerability information to harm non-subscribers [30].

Several variations on vulnerability markets have been proposed. Böhme has argued that software derivatives might be better [31]. Contracts for software would be issued in pairs: the first pays a fixed value if no vulnerability is found in a program by a specific date, and the second pays another value if one is found. If these contracts can be traded, then their price should reflect the consensus on software quality. Software vendors, software company investors, and insurance companies could use such derivatives to hedge risks. A third possibility, due to Ozment, is to design a vulnerability market as an auction; this has recently been implemented by the Swiss firm WabiSabiLabi [32].

One criticism of all market-based approaches is that they might increase the number of identified vulnerabilities by motivating more people to search flaws. Thus some care must be exercised in designing them.

An alternative approach is insurance. Underwriters often use expert assessors to look at a client firm's IT infrastructure and management; this provides data to both the insured and the insurer. Over the long run, insurers learn to value risks more accurately. Right now, however, the cyber-insurance market is both underdeveloped and underutilised. One reason, according to Böhme and Kataria [33], is the interdependence of risk, which takes both local and global forms. Firms' IT infrastructure is connected to other entities – so their efforts may be undermined by failures elsewhere. Cyber-attacks often exploit a vulnerability in a program used by many firms. Interdependence can make some cyber-risks unattractive to insurers – particularly those risks that are globally rather than locally correlated, such as worm and virus attacks, and systemic risks such as Y2K.

Many writers have called for software risks to be transferred to the vendors; and, despite the disclaimers in software end user license agreements (EULAs), the consumer protection laws in both the USA and the EU allow people to sue for personal injury and in some cases damage following software failure. But this is still not common. (We'll return to this later.) If it becomes so, then it could be interesting to see whether large platform vendors such as Microsoft are able to buy insurance. Even at the level of customer firms, correlated risk makes firms under-invest in both security technology and cyber-

insurance [34]. Cyber-insurance markets may in any case lack the volume and liquidity to become efficient.

3.2 Economics of privacy

The persistent erosion of personal privacy has frustrated policy makers and practitioners alike. People say that they value privacy, yet act otherwise. Privacy-enhancing technologies have been offered for sale, yet most have failed in the marketplace. Why should this be?

Privacy is one aspect of information security that interested economists before 2000. In 1978, Posner defined privacy in terms of secrecy [35], and the following year extended this to seclusion [36]. In 1980, Hirshleifer published a seminal paper in which he argued that rather than being about withdrawing from society, privacy was a means of organising society, arising from evolved territorial behaviour; internalised respect for property is what allows autonomy to persist in society. These privacy debates in the 1970s led in Europe to generic data-protection laws, while the USA limited itself to a few sector-specific laws such as HIPAA. Economists' appetite for work on privacy was further whetted recently by the Internet, the dotcom boom, and the exploding trade in personal information about online shoppers.

An early modern view of privacy can be found in a 1996 paper by Varian who analysed privacy in terms of information markets [38]. Consumers want to not be annoyed by irrelevant marketing calls while marketers do not want to waste effort. Yet both are frustrated, because of search costs, externalities and other factors. Varian suggested giving consumers rights in information about themselves, and letting them lease it to marketers with the proviso that it not be resold without permission.

The recent proliferation of complex, information-intensive business models demand a broader approach. Odlyzko argued in 2003 that privacy erosion is a consequence of the desire to charge different prices for similar services [39]. Technology is simultaneously increasing both the incentives and the opportunities for price discrimination. Companies can mine online purchases and interactions for data revealing individuals' willingness to pay. From airline yield-management systems to complex and ever-changing software and telecommunications prices, differential pricing is economically efficient – but increasingly resented. Acquisti and Varian analyzed the market conditions under which personalised price discrimination is profitable [40]: it may thrive in industries with wide variation in consumer valuation for services, where services can be personalised at low marginal cost, and where repeated purchases are likely.

Acquisti and Grossklags tackled the specific problem of why people express a high preference for privacy when interviewed but reveal a much lower preference through their behaviour both online and offline [41]. They find that people mostly lack sufficient information to make informed choices, and even when they do they often trade long-term privacy for short-term benefits. Loewenstein developed this into a fuller psychology-and-economics analysis (a topic to which we'll return later) [42]. By combining this 'hyperbolic discounting' with loss aversion, adaptation and preference uncertainty, he predicts much of the oberved behaviour: that people will initially oppose any loss of privacy but will rapidly adapt and will not be very motivated to gain new forms of privacy. Thus unless there is some crisis at which a lot of privacy looks like being lost at one time, people may

get used gradually to levels of surveillance that would have been unacceptable if imposed quickly. Yet new technologies magnify the risks while removing the cues; so we may eventually need either more regulation or greater tolerance for individual idiosyncrasies.

Swire argued that we should measure the costs of privacy intrusion more broadly [44]. If a telesales operator calls 100 prospects, sells three of them insurance, and annoys 80, then the conventional analysis considers only the benefit to the three and to the insurer. However, persistent annoyance causes millions of people to go ex-directory, to not answer the phone during dinner, or to screen calls through an answering machine. The long-run societal harm can be considerable. Several empirical studies have backed this up by examining people's privacy valuations. Vila et al. further characterised privacy economics as a lemons market [43], arguing that consumers disregard future price discrimination when giving information to merchants.

So much for the factors that make privacy intrusions more likely. What factors make them less so? Campbell et al. found that the stock price of companies reporting a security breach is more likely to fall if the breach leaked confidential information [45]. Acquisti, Friedman and Telang conducted a similar analysis for privacy breaches [46]. Their initial results are less conclusive but still point to a negative impact on stock price followed by an eventual recovery.

Regulatory responses (pioneered in Europe) have largely centred on requiring companies to allow consumers to either 'opt-in' or 'opt-out' of data collection. While privacy advocates typically support opt-in policies as they result in lower rates of data collection, Bouckaert and Degryse argue for opt-out on competition grounds [47]: the availability of information about the buying habits of most customers, rather than a few customers, may help competitors to enter a market.

Empirically, there is wide variation in 'opt-out' rates between different types of consumer, but their motives are not always clear. Varian et al. analyzed the FCC's telephone-sales blacklist by district [48]. They found that educated people are more likely to sign up: but is that because rich households get more calls, because they value their time more, or because they understand the risks better?

Incentives also affect the design of privacy technology. Builders of anonymity systems know they depend on network externalities: more users mean more cover traffic to hide activities from the enemy [49]. An interesting case is Tor [50], which anonymises web traffic and emphasises usability to increase adoption rates. It developed from a US Navy communications system, but eventually all internet users were invited to participate in order to build network size, and it is now the largest anonymous communication system known.

3.3 Incentives and the deployment of security mechanisms

Insurance is not the only market affected by information security. Some very high-profile debates have centred on DRM; record companies have pushed for years for DRM to be incorporated into computers and consumer electronics, while digital-rights activists have opposed them. What light can security economics shed on this debate?

Many researchers have set the debate in a much wider context than just record companies versus downloaders. Varian pointed out in 2002 that DRM and similar mechanisms were

also about tying, bundling and price discrimination; and that their unfettered use could damage competition [51]. A paper by Samuelson and Scotchmer studied what might go wrong if technical and legal restraints were to undermine the right to reverse engineer software products for compatibility. It provided the scholarly underpinnings for much of the work on the anti-competitive effects of the DMCA, copyright control mechanisms, and information security mechanisms applied to new business models.

'Trusted Computing' (TC) mechanisms have come in for significant analysis and criticism. Von Hippel showed how most of the innovations that spur economic growth are not anticipated by the manufacturers of the platforms on which they are based; the PC, for example, was conceived as an engine for running spreadsheets, and if IBM had been able to limit it to doing that, a huge opportunity would have been lost. Furthermore, technological change in IT markets is usually cumulative. If security technology can be abused by incumbent firms to make life harder for innovators, this will create all sorts of traps and perverse incentives [53]. Anderson pointed out the potential for competitive abuse of the TC mechanisms; for example, by transferring control of user data from the owner of the machine on which it is stored to the creator of the file in which it is stored, the potential for lock-in is hugely increased [27]. Lookabaugh and Sicker discussed an existing case history of an industry crippled by security-related technical lock-in [54]. US cable industry operators are locked in to their set-top-box vendors; and although they largely negotiated away the direct costs of this when choosing a suppler, the indirect costs were large and unmanageable. Innovation suffered and cable fell behind other platforms, such as the Internet, as the two platform vendors did not individually have the incentive to invest in improving their platforms.

Economic research has been applied to the record industry itself, with results it found disturbing. In 2004, Oberholzer and Strumpf published a now-famous paper, in which they examined how music downloads and record sales were correlated [55]. They showed that downloads do not do significant harm to the music industry. Even in the most pessimistic interpretation, five thousand downloads are needed to displace a single album sale, while high-selling albums actually benefit from file sharing.

In January 2005, Varian presented a surprising result [56]: that stronger DRM would help system vendors more than the music industry, because the computer industry is more concentrated (with only three serious suppliers of DRM platforms – Microsoft, Sony, and the dominant firm, Apple). The content industry scoffed, but by the end of that year music publishers were protesting that Apple was getting too large a share of the cash from online music sales. As power in the supply chain moved from the music majors to the platform vendors, so power in the music industry appears to be shifting from the majors to the independents, just as airline deregulation favoured aircraft makers and low-cost airlines. This is a striking demonstration of the predictive power of economic analysis. By fighting a non-existent threat, the record industry had helped the computer industry forge a weapon that may be its undoing.

3.4 Protecting computer systems from rational adversaries

Information security practitioners traditionally assumed two types of user: honest ones who always behave as directed, and malicious ones intent on wreaking havoc at any cost. But systems are often undermined by what economists call *strategic* users: users who

act out of self-interest rather than malice. Many file-sharing systems suffer from 'free-riding', where users download files without uploading their own. This is perfectly rational behaviour, given that upload bandwidth is typically more scarce and file uploaders are at higher risk of getting sued. The cumulative effect is degraded performance.

Another nuisance caused by selfish users is spam. The cost per transmission to the spammer is so low that a tiny success rate is acceptable [57]. Furthermore, while spam imposes significant costs on recipients, these costs are not felt by the spammers. Böhme and Holz examined stock spam and identified statistically significant increases in the price of touted stocks [58]. Frieder and Zittrain independently find a similar effect [59].

Several network protocols may be exploited by selfish users at the expense of system-wide performance. In TCP, the protocol used to transmit most Internet data, Akella et al. find that selfish provision of congestion control mechanisms can lead to suboptimal performance [60].

Researchers have used game theory to study the negative effects of selfish behaviour on systems more generally. Koutsoupias and Papadimitriou termed the 'price of anarchy' as the ratio of the utilities of the worst-case Nash equilibrium to the social optimum [61]. The price of anarchy has become a standard measurement of the inefficiency of selfish behaviour in computer networks. Roughgarden and Tardos studied selfish routing in a congested network, comparing congestion levels in a network where users choose the shortest path available to congestion when a network planner chooses paths to maximise flow [62]. They established an upper bound of $\frac{4}{3}$ for the price of anarchy when congestion costs are linear; furthermore, in general, the total latency of a selfish network is at most the same as an optimal flow routing twice as much traffic.

Other topics hindered by selfish activity include network creation, where users decide whether to create costly links to shorten paths or free-ride over longer, indirect connections [63, 64, 65]; wireless spectrum sharing, where service providers compete to acquire channels from access points [66]; and computer virus inoculation, where users incur a high cost for inoculating themselves and the benefits accrue to unprotected nodes [67].

To account for user self-interest, computer scientists have proposed several mechanisms with an informal notion of 'fairness' in mind. To address spam, Dwork and Naor propose attaching to emails a 'proof-of-work' that is easy to do for a few emails but impractical for a flood [68]. Laurie and Clayton criticise 'proof-of-work' schemes, demonstrating that the additional burden may be cumbersome for many legitimate users while spam senders could use botnets to perform the computations [69]. Furthermore, ISPs may not be prepared to block traffic from these compromised machines. Serjantov and Clayton analyse the incentives on ISPs to block traffic from other ISPs with many infected machines, and back this up with data [70]. They also show how a number of existing spam-blocking strategies are irrational and counterproductive.

Reputation systems have been widely proposed to overcome free-riding in peer-to-peer networks. The best-known fielded example may be feedback on eBay's online auctions. Dellarocas argues that leniency in the feedback mechanism (only 1% of ratings are negative) encourages stability in the marketplace [72]. Serjantov and Anderson use social choice theory to recommend improvements to reputation system proposals [73]. Feldman et al model such systems as an iterated prisoner's dilemma game, where users in each round alternate between roles as client and server [71]. Recently, researchers have begun

to consider more formally how to construct fair systems using mechanism design. We discuss these developments in Section 5.1.

4 The Role of Governments

The information security world has been regulated from the beginning, although initially government concerns had nothing to do with competition policy. The first driver was a non-proliferation concern. Governments used export licenses and manipulated research funding to restrict access to cryptography for as long as possible. This effort was largely abandoned in 2000. The second driver was the difficulty that even the US government had over many years in procuring systems for its own use, once information security came to encompass software security too. Thus, during the 80s and 90s, it was policy to promote research in security while hindering research in cryptography.

4.1 Early days

Landwehr describes the efforts of the US government from the mid-1980s to tackle a the lemons problem in the security software business [74]. The first attempted fix was a government evaluation scheme – the Orange Book – but that brought its own problems. Managers' desire for the latest software eased certification requirements: vendors had to simply show that they had initiated the certification process, which often was never completed. Evaluations were also conducted at government expense by NSA civil servants, who being risk-averse took their time; evaluated products were often unusably out of date. There were also problems interworking with allies' systems, as countries such as the UK and Germany had their own incompatible schemes.

This led the NATO governments to establish the 'Common Criteria' as a successor to the Orange Book. Most evaluations are carried out by commercial laboratories and are paid for by the vendor who is supposed to be motivated by the cachet of a successful evaluation. The Common Criteria suffer from different problems, most notably adverse selection: vendors shop around for the evaluator who will give them the easiest ride, and the national agencies who certify the evaluation labs are very reluctant to revoke a license, even following scandal, because of fears that confidence in the scheme will be undermined [75].

Regulation is increasingly justified by perceived market failures in the information security industry. The European Union has proposed a Network Security Policy that sets out a common European response to attacks on information systems [76]. This starts using economic arguments about market failure to justify government action in this sector. The proposed solutions are familiar, involving everything from consciousness raising to more Common Criteria evaluations.

Another explicit use of security economics in policymaking was the German government's comments on Trusted Computing [77]. These set out concerns about issues from certification and trapdoors through data protection to economic policy matters. They were hugely influential in persuading the Trusted Computing Group to incorporate and adopt membership rules that mitigated the risk of its program discriminating against small-to-medium sized enterprises. Recently the European Commission's DG Competition has been considering the economic implications of the security mechanisms of Vista.

Among academic scholars of regulation, Barnes studies the incentives facing the virus writers, software vendors and computer users [78], and contemplates various policy initiatives to make computers less liable to infection, from rewarding those who discover vulnerabilities to penalising users who do not adopt minimal security standards. Garcia and Horowitz observe that the gap between the social value of internet service providers, and the revenue at stake associated with their insecurity, is continuing to increase [79]. If this continues, they argue, mandatory security standards may become likely.

Moore presents an interesting regulatory question from forensics. While PCs use standard disc formats, mobile phones use proprietary interfaces, which make data recovery from handsets difficult; recovery tools exist only for the most common models. So criminals should buy unfashionable phones, while the police should push for open standards [80].

4.2 Self-regulation

Heavy-handed regulation can introduce high costs – whether directly, or as a result of agency issues and other secondary factors. Ghose and Rajan discuss how three US laws – Sarbanes-Oxley, Gramm-Leach-Bliley and HIPAA – place a disproportionate burden on small and medium sized businesses, largely through a one-model-fits-all approach to compliance by the big accounting firms [81]. They show how mandatory investment in security compliance can create unintended consequences from distorting security markets to reducing competition.

Given the high costs and doubtful effectiveness of regulation, self-regulation has been tried in a number of contexts, but some attempts failed spectacularly. For example, a number of organisations have set up certification services to vouch for the quality of software products or web sites. Their aim was twofold: to overcome public wariness about electronic commerce, and to forestall more expensive regulation by the government. But (as with the Common Criteria) certification markets can easily be ruined by a race to the bottom; dubious companies are more likely to buy certificates than reputable ones, and even ordinary companies may shop around for the easiest deal. In the absence of a capable motivated regulator, ruin can arrive quickly.

Edelman analysed this 'adverse selection' in the case of website approvals and online advertising [82]: while about 3% of websites are malicious, some 8% of websites with certification from one large vendor are malicious. He also compared ordinary web search results and those from paid advertising, finding that while 2.73% of companies ranked top in a web search were bad, 4.44% of companies who had bought ads from the search engine were bad. His conclusion – 'Don't click on ads' – could be bad news for the search industry.

Self-regulation has fared somewhat better for patch management. Analysis by Arora et al. shows that competition in software markets hastens patch release even more than the threat of vulnerability disclosure in two out of three studied strategies [84]. Beattie et al. found that pioneers who apply patches quickly end up discovering problems that break their systems, but laggards are more vulnerable to attack [83].

Governments also facilitate the sharing of security information between private companies. Two papers analyse the incentives that firms have to share information on security breaches within the Information Sharing and Analysis Centers (ISACs) set up after 9/11

by the US government [85, 86]. Theoretical tools developed to model trade associations and research joint ventures can be applied to work out optimal membership fees and other incentives.

4.3 Practical policy options

A report commissioned by the European Network and Information Security Agency (and of which we are among the authors) explored practical policy actions to align incentives for improving information security at the EU level [87]. There has long been a shortage of hard data about information security failures, as many of the available statistics are not only poor but are collected by parties such as security vendors or law enforcement agencies that have a vested interest in under- or over-reporting. Banks and ISPs are two particularly problematic 'black holes' where data are fragmentary or simply unavailable. Several recommendations aim to collect and publish robust statistics, beginning with the adoption of comprehensive security-breach disclosure laws and the collection of consistent Europe-wide fraud statistics.

It is widely known that well-run ISPs are diligent about identifying and quarantining infected machines, while badly-run ISPs are not. The first step in tackling this problem is to collect better data about the quantity of spam and other bad traffic emitted by ISPs. However, transparency alone isn't enough, since infected machines are often not harming their owners but are instead polluting the digital environment by sending spam, hosting phishing websites and distributing illegal content.

These externalities must be internalised, and the report recommends a fairly radical approach: introducing a a statutory scale of damages against ISPs that do not respond promptly to requests for the removal of compromised machines. Fixed penalties are useful because they avoid the problem of quantifying losses following every infringement. They have been used effectively in the airline industry, where the EU has introduced penalties for airlines that deny passengers boarding due to overbooking, cancellations or excessive delays. The goal of fixed penalties is to provide an effective deterrent, while simplifying the liability when violations occur.

Another highlight of the report is its recommendations on liability for defective software. It takes the pragmatic view that software liability is too large an issue to be dealt with in one go, because of the large and growing variety of goods and services in which software plays a critical role. It is also important not to discriminate between different service modalities. For example, a navigation service can be delivered as a GPS device, an online service, software for a phone or PDA or laptop, or some combination of these; it would distort the market if a lorry driver who relied on a defective service and got stuck in a narrow lane could only sue if the service had been delivered in some particular way.

The report recommends a staged approach that would leave standalone products and services to be dealt with by existing regulations on safety, product liability and consumer rights. Europe already has many regulations on standards for everything from cars to TV sets. The new issue is connectivity. What if network-connected devices become infected, and start distributing spam or phish? As networked systems can cause harm to others, it is proposed that Europe develop appropriate standards and require vendors to certify that their products are secure by default. If the certification later turns out to have been wrong, the vendor would become liable and ISPs could recover penalty charges from it.

5 Open Problems

There are many active areas of security-economics research. Here we highlight just four live problems. Each lies not just at the boundary between security and economics, but also at the boundary between economics and some other discipline – respectively algorithmic mechanism design, network science, organisational theory and psychology.

5.1 Algorithmic Mechanism Design

Given the largely unsatisfactory impact of information security regulation, a complementary approach based on mechanism design is emerging. Researchers are beginning to design network protocols and interfaces that are 'strategy-proof': that is, designed so that no-one can gain by cheating [88]. Designing bad behavior out of systems may be cheaper than policing it afterwards.

One key challenge is to allocate scare digital resources fairly, and the theory of auctions may provide the key. Nisan and Segal show that although one can solve the allocation problem using strategy-proof mechanisms, the number of bits that must be communicated grows exponentially; thus in many cases the best practical mechanism will be a simple bundled auction [89]. They also suggest that if arbitrary valuations are allowed, players can submit bids that will cause communications complexity problems for all but the smallest auctions.

Some promising initial results look at mechanism design and protocols. Feigenbaum et al. show how combinatorial auction techniques can be used to provide distributed strategy-proof routing mechanisms [90]. Schneidman et al. compare the incentive mechanisms in BitTorrent, a popular peer-to-peer file-sharing application, to theoretical guarantees of faithfulness [91].

5.2 Network topology and information security

There has been an interesting collaboration recently between physicists and sociologists in analyzing the topology of complex networks and its effect on social interactions. Computer networks, like social networks, are complex but emerge from ad-hoc interactions of many entities using simple ground rules. The new discipline of network analysis takes ideas from sociology, condensed-matter physics and graph theory, and in turn provides tools for modelling and investigating such networks (see [92] for a recent survey). Some economists have also recognised the impact of network structure on a range of activities, from crime [93, 94] to the diffusion of new technologies [95]. Other researchers have focused on why networks are formed, where the individual costs of establishing links between agents is weighed against the overall benefit of improved connectivity [96]. Economic models are well-suited to comparing the social efficiency of different network types and predicting which structures are likely to emerge when agents act selfishly. See [97] for a collection of recent work.

Network topology can strongly influence conflict dynamics. Often an attacker tries to disconnect a network or increase its diameter by destroying nodes or edges, while the defender counters using various resilience mechanisms. Examples include a music industry body attempting to close down a peer-to-peer file-sharing network; a police force trying to decapitate a terrorist organisation; and a totalitarian government harrassing political

activists. Police forces have been curious for some years about whether network science might be of practical use in covert conflicts – whether to insurgents or to counterinsurgency forces.

Different topologies have different robustness properties. Albert, Jeong and Barabási showed that certain real world networks with scale-free degree distributions resist random attacks much better than targeted attacks [98]. This is because scale-free networks – like many real-world networks – get much of their connectivity from a few nodes with high vertex order. This resilience makes them highly robust against random upsets; but remove the 'kingpin' nodes, and connectivity collapses.

This is the static case – for example, when a police force becomes aware of a criminal or terrorist network, and sets out to disrupt it by finding and arresting its key people. Nagaraja and Anderson extend this to the dynamic case. In their model, the attacker can remove a certain number of nodes at each round, after which the defenders recruit other nodes to replace them [99]. They studied how attack and defence interact using multi-round simulations, and found that forming localised clique structures at key network points works reasonably well while defences based on rings did not work well at all. This helps explain why peer-to-peer systems with ring architectures turned out to be rather fragile – and why revolutionaries have tended to organise themselves in cells.

An open challenge is how to reconcile the differences between generated network models and computer networks. Degree distribution is only one factor in the structure of a network. Li et al. closely examined the topology of computer networks [100] and found that degree-centrality attacks on the Internet do not work well since edge routers that connect to homes have much higher degree than backbone routers at major IPSs. For attacks on privacy, however, topological analysis has proven quite effective. When Danezis and Wittneben applied these network analysis ideas to privacy [101], they found that doing traffic analysis against just a few well-connected organisers can draw a surprising number of members of a dissident organisation into the surveillance net.

5.3 Large project management

As well as extending into system design, crime, and covert conflict, security economics may help the student of information systems management. Perhaps the largest issue here is the risk of large software project failures, which can cost billions and threaten the survival of organisations.

We noted above that perhaps 30% of large development projects fail [24], and this figure seems impervious to technological progress: better tools help engineers make larger systems, the same proportion of which still fail as before. This suggests that project failure is not technical but down to socio-economic factors such as the way decisions are taken in firms. There is thus a temptation to place what we now know about the economics of dependability alongside institutional economics and perform a gap analysis.

One interesting question is whether public-sector organisations are particularly prone to large software project failure. The CIO of the UK's Department of Work and Pensions recently admitted that only 30% of government IT projects succeed [102]. There are many possible reasons. The dependability literature teaches that large software project failures are mostly due to overambitious, vague or changing specifications, coupled with

poor communications and an inability to acknowledge the signs of failure early enough to take corrective action. Good industrial project managers try to close down options fast, and get the customer to take the hard decisions upfront. Elected politicians, on the other hand, are in the business of mediating conflicts between different interests and groups in society, and as many of these conflicts are transient, avoiding or delaying hard choices is a virtue. Furthermore, at equilibrium, systems have too many features because the marginal benefit of the typical feature accrues to a small vocal group, while the cost is distributed across a large user base as a slightly increased risk of failure. This equilibrium may be even further from the optimum when design decisions are taken by elected officials: the well-known incentives to dump liability, to discount consequences that will arrive after the next election or reshuffle, and to avoid ever admitting error, surely add their share. The economics of dependability may thus be an interesting topic for researchers in schools of government.

5.4 Psychology and security

The final point at which security economics is sparking off another discipline is with psychology – a field with which security engineers have already had at least three points of contact. First, three famous experiments in social psychology showed the ease with which people could be bullied by authority figures, or persuaded by peers, to behave inappropriately. In 1951, Solomon Asch showed that most people could be induced to deny the evidence of their own eyes in order to conform to a group [103]; in 1961, Milgram showed that most people would administer severe electric shocks to an actor playing the role of a 'learner' at the behest of an experimenter playing the role of the 'teacher' – even when the 'learner' appeared to be in severe pain and begged the subject to stop [104]; and in 1971, the Stanford Prisoner Experiment showed that normal people can egg each other on to behave wickedly even in the absence of orders. There, students playing the role of warders so brutalised students playing the role of prisoners that the experiment had to be stopped [105].

Inappropriate obedience is a live problem: card thieves call up cardholders, pretend to be from the bank, and demand the PIN [2, 75]. Worse, in 1995-2005, a hoaxer calling himself 'Officer Scott' ordered the managers of dozens of US stores and restaurants to detain some young employee on suspicion of theft and strip-search her or him. Various other degradations were ordered, including beatings and sexual assaults. At least 13 people who obeyed the caller and did searches were charged with crimes, and seven were convicted [106]. In October 2007, a jury ordered McDonalds to pay $6.1 million dollars to one of the victims, who had been strip-searched and indecently assaulted when an 18-year-old employee [107].

The second point of contact has been security usability, which has become a growth area recently; early results are collected in [108]. The third has been the study of deception – a somewhat less well-defined field, but which extends from conjuring to camouflage to the study of fraud, and which is interesting the security usability community more as phishing becomes a serious problem.

There is a potentially valuable interface with economics here too. Economic analysis traditionally assumed that the principals are rational and act out of pure self-interest. Real people depart in a number of ways from this ideal, and there has arisen in recent

years a vigorous school of economic psychology or behavioural economics, which studies the effects that human social and cognitive biases have on economic decision-making. The 2002 Nobel prize was awarded to Daniel Kahnemann for his seminal role in establishing this field, and particularly in decision-making under risk and uncertainty. We already referred to behavioral economics in the context of privacy, and its extensive studies of the heuristics and biases we use in our daily lives have much more to teach the security engineer [109]. In addition to loss aversion, time discounting and preference uncertainty, there are issues of *framing* – an action as a gain rather than as a loss makes people more likely to take it; the *availability heuristic* whereby easily-remembered data have more weight in mental processing; the *anchoring effect* whereby we base a judgement on an initial guess or comparison and then adjust it if need be; channel biases (we're more likely to be sceptical about things we've heard than about things we've seen) and biases in the way we deal with mental accounting, ranging from reluctance to write off money that we have wasted to a tendency to worry too much about unlikely events.

Schneier has discussed cognitive biases as the root cause of our societies' vulnerability to terrorism [110]. The psychologist Daniel Gilbert, in an article provocatively entitled 'If only gay sex caused global warming', also discusses why we are much more afraid of terrorism than of climate change [111]. We have many built-in biases that made perfect evolutionary sense when we were living in small social groups on the plains of Africa half a million years ago, but may now be maladaptive. For example, we are more sensitive to risks involving intentionality, whether of a person or animal, as the common causes of violent death back then included hungry lions and enemies with sharp sticks. We are also more afraid of uncertainty; of rare or unfamiliar risks; of risks controlled by others, particularly 'outsiders' or other people we don't trust or find morally offensive. A number of these biases tie in with defects in our mental accounting.

The study of cognitive biases may also help illuminate fraud and phishing. The *fundamental attribution error* – that people often err by trying to explain things by intentionality when their causes are in fact impersonal – undermines efforts to curb phishing by teaching users about the gory design details of the Internet – for example, by telling them to parse URLs in emails that seem to come from a bank. As soon as users get confused, they will revent to judging a website by its 'look and feel'.

One potential area of research is gender. Recently people have realised that software can create barriers to females, and this has led to research work on 'gender HCI' – on how software should be designed so that women as well as men can use it effectively. The psychologist Simon Baron-Cohen classifies human brains into type S (systematizers) and type E (empathizers) [112]. Type S people are better at geometry and some kinds of symbolic reasoning, while type Es are better at language and multiprocessing. Most men are type S, while most women are type E. Of course, innate abilities can be modulated by many developmental and social factors. Yet, even at a casual reading, this material raises a suspicion that many security mechanisms are far from gender-neutral. Is it unlawful sex discrimination for a bank to expect its customers to detect phishing attacks by parsing URLs?

Yet another bundle of issues concern how we evolved our perceptions of ourselves, and – critically – of others. Another interesting insight from Baron-Cohen's work is that humans are most distinct from other primates in that we have a theory of mind; our

brains are wired so that we can imagine others as being like ourselves, to empathise with them better. Indeed there is a thread of research in psychology and primatology going back to the 1970s which argues that we developed our intelligence in a social context: the positive way of putting this is that coping with complex social groups became adaptive, while the more cynical version is that people who were good at deception, or at detecting deception in others, had more surviving offspring. No doubt our capabilities co-evolved over many generations of lies, social manipulation, sexual infidelities and revenge. This 'Machiavellian brain' hypothesis, developed by Humphrey, Byrne, Whiten and others [113] must be tempered by the observation that as the more social (type E) individuals tend to be women, the social brain may have been at least as much tied up with nurturing the young. Then again, archaeologists such as Steven LeBlanc have shown that most humans lived at most times in a condition of chronic tribal warfare – as indeed bands of chimps do to this day [114]. The resulting evolutionary pressures have equipped us to cooperate with members of our own group, but also to see members of competing groups as predators, prey or even vermin, and deal with them cruelly [115]. Such research is clearly of relevance to the security engineer, as well as fascinating: our trade has been tied up for thousands and perhaps millions of years with human intellectual and cultural development.

6 Conclusions

Over the last few years, a research program on the economics of security has built many cross-disciplinary links and has produced many useful (and indeed delightful) insights from unexpected places. Many perverse things, long known to security practitioners but just dismissed as 'bad weather', turn out to be quite explicable in terms of the incentives facing individuals and organisations, and in terms of different kinds of market failure.

As for the future, the work of the hundred or so researchers active in this field has started to spill over into at least four new domains. The first is the technical question of how we can design better systems by making protocols strategy-proof so that the incentives for strategic or malicious behaviour are removed a priori.

The second is the economics of security generally, where there is convergence with economists studying topics such as crime and warfare. The causes of insurgency, and tools for understanding and dealing with insurgent networks, are an obvious attractor.

The third is the economics of dependability. Large system failures cost industry billions, and the problems seem even more intractable in the public sector. We need a better understanding of what sort of institutions can best evolve and manage large complex interconnected systems.

Finally, the border between economics and psychology seems particularly fruitful, both as a source of practical ideas for designing more usable secure systems, and as a source of deeper insights into foundational issues.

Acknowledgments This is an evolving review paper of a rapdily-developing field. An early version of it was presented at SoftInt 2007, and a later version at Crypto 2007. The authors have also given a number of invited talks on the subject, and are grateful for

much feedback from audiences and from colleagues. Tyler Moore is supported by the UK Marshall Aid Commemoration Commission and the US National Science Foundation.

Refernces

[1] Michael Mastanduno, "Economics and Security in Statecraft and Scholarship", *International Organization* v 52 no 4 (Autumn 1998)

[2] Ross Anderson, "Why Cryptosystems Fail", in *Communications of the ACM* v 37 no 11 (Nov 94) pp 32–40

[3] Ian Ayres, Steven Levitt, "Measuring Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack", NBER Workign Paper no W5928; also in *The Quarterly Journal of Economics* v 113 pp 43–77

[4] Jean Camp, Catherine Wolfram, "Pricing Security", in *Proceedings of the CERT Information Survivability Workshop* (Oct 24-26 2000) pp 31–39

[5] Hal Varian, Managing Online Security Risks, Economic Science Column, The New York Times, June 1, 2000

[6] Nick Bohm, Ian Brown and Brian Gladman, "Electronic Commerce: Who Carries the Risk of Fraud?" in *Journal of Information, Law and Technology* v 3 York Times, June 1, 2000

[7] Ross Anderson, "Closing the Phishing Hole – Fraud, Risk and Nonbanks", in *Nonbanks in the Payment System*, Santa Fe, May 2007

[8] Tyler Moore, "Countering Hidden-Action Attacks on Networked Systems", in *Fourth Workshop on the Economics of Information Security*, June 2005, Harvard.

[9] Ross Anderson, "The Eternity Service", in *Pragocrypt 96*

[10] George Danezis and Ross Anderson, "The Economics of Resisting Censorship", in *IEEE Security & Privacy* v 3 no 1 (2005) pp 45–50

[11] David Goodhart, "Too Diverse?", in *Prospect* (Feb 2004) and at `http://www.guardian.co.uk/race/story/0,11374,1154684,00.html`

[12] Ross J. Anderson, "Why Information Security is Hard – An Economic Perspective", in *17th Annual Computer Security Applications Conference* (Dec 2001) and at `http://www.cl.cam.ac.uk/users/rja14/Papers/econ.pdf`

[13] Jack Hirshleifer, "From weakest-link to best-shot: the voluntary provision of public goods", in *Public Choice* v 41, (1983) pp 371–386

[14] Hal Varian, "System Reliability and Free Riding", in *Economics of Information Security*, Kluwer 2004 pp 1–15

[15] Howard Kunreuther and Geoffrey Heal, "Interdependent Security", in *Journal of Risk and Uncertainty* v 26 no 2–3 (March-May 2003) pp 231–249

[16] Michael Katz and Carl Shapiro, "Network Externalities, Competition, and Compatibility", in *The American Economic Review* v 75 no 3 (June 1985) pp 424–440

[17] Andy Ozment and Stuart Schechter, "Bootstrapping the Adoption of Internet Security Protocols", Fifth Workshop on the Economics of Information Security (June 26–28, Cambridge, UK)

[18] Ross Anderson, "Open and Closed Systems are Equivalent (that is, in an ideal world)", in *Perspectives on Free and Open Source Software*, MIT Press 2005, pp 127–142

[19] Eric Rescorla, "Is Finding Security Holes a Good Idea?", Third Workshop on the Economics of Information Security (2004)

[20] Andy Ozment, "The Likelihood of Vulnerability Rediscovery and the Social Utility of Vulnerability Hunting", Fourth Workshop on the Economics of Information Security (2005)

[21] Andy Ozment and Stuart Schechter, "Milk or Wine: Does Software Security Improve with Age?" in *15th Usenix Security Symposium* (2006)

[22] Ashish Arora, Rahul Telang and Hao Xu, "Optimal Policy for Software Vulnerability Disclosure", Third Workshop on the Economics of Information Security (May 2004, Minneapolis, MN)

[23] Ashish Arora, Ramayya Krishnan, Anand Nandkumar Rahul Telang and Yubao Yang, Impact of Vulnerability Disclosure and Patch Availability – An Empirical Analysis, Third Workshop on the Economics of Information Security (2004)

[24] Bill Curtis, Herb Krasner, Neil Iscoe, "A Field Study of the Software Design Process for Large Systems", in *Communications of the ACM* v 31 no 11 (Nov 88) pp 1268–1287

[25] Carl Shapiro, Hal Varian, *'Information Rules'*, Harvard Business School Press (1998)

[26] George Akerlof, "The Market for 'Lemons: Quality Uncertainty and the Market Mechanism", in *The Quarterly Journal of Economics* v 84 no 3 (1970) pp 488–500

[27] Ross Anderson, "Cryptography and Competition Policy – Issues with 'Trusted Computing' ", Second Workshop on Economics and Information Security (2003)

[28] VISA, *PIN Management Requirements: PIN Entry Device Security Requirements Manual* (2004)

[29] Stuart Schechter, "Computer Security Strength & Risk: A Quantitative Approach", Harvard University, May 2004

[30] Karthik Kannan and Rahul Telang, "Economic Analysis of Market for Software Vulnerabilities", Third Workshop on the Economics of Information Security (2004)

[31] Rainer Böhme, "A Comparison of Market Approaches to Software Vulnerability Disclosure", in *ETRICS 2006* Springer LNCS v 2995 pp 298–311

[32] Andy Ozment, "Bug Auctions: Vulnerability Markets Reconsidered", Third Workshop on the Economics of Information Security (2004)

[33] Rainer Böhme and Gaurav Kataria, "Models and Measures for Correlation in Cyber-Insurance", Fifth Workshop on the Economics of Information Security (2006)

[34] Hulisi Ogut, Nirup Menon and Srinivasan Raghunathan, "Cyber Insurance and IT Security Investment: Impact of Interdependent Risk", in Fourth Workshop on the Economics of Information Security (2005)

[35] Richard Posner, "An Economic Theory of Privacy", in *Regulation* (1978) pp 19–26

[36] Richard Posner, "Privacy, Secrecy and Reputation" in *Buffalo Law Review* v 28 no 1 (1979)

[37] Jack Hirshleifer, "Privacy: its Origin, Function and Future", in *Journal of Legal Studies* v 9 (Dec 1980) pp 649–664

[38] Hal Varian, "Economic Apects of Personal Privacy", in *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration report, 1996

[39] Andrew Odlyzko, "Privacy, economics, and price discrimination on the Internet", in *ICEC '03: Proceedings of the 5th international conference on Electronic commerce* pp 355–366

[40] Alessandro Acquisti and Hal Varian, "Conditioning Prices on Purchase History" in *Marketing Science* v 24 no 3 (2005)

[41] Alessandro Acquisti and Jens Grossklags, "Privacy and Rationality: Preliminary Evidence from Pilot Data", in Third Workshop on the Economics of Information Security (2004, Minneapolis, Mn)

[42] G Loewenstein, "Searching for Privacy in All the Wrong Places: A behavioral economics perspective on individual concern for privacy", at Workshop on the Economics of Information Security 2007

[43] Tony Vila, Rachel Greenstadt and David Molnar, "Why we can't be bothered to read privacy policies", in *Economics of Information Security* (Kluwer, 2004) pp 143–154

[44] P Swire, "Efficient Confidentiality for Privacy, Security, and Confidential Business Information", Brookings-Wharton Papers on Financial Services (Brookings, 2003)

[45] Katherine Campbell, Lawrence Gordon, Martin Loeb and Lei Zhou, "The economic cost of publicly announced information security breaches: empirical evidence from the stock market", in *Journal of Computer Security* v 11 no 3 (2003) pp 431–448

[46] Alessandro Acquisti, Allan Friedman and Rahul Telang, "Is There a Cost to Privacy Breaches?", Fifth Workshop on the Economics of Information Security (2006)

[47] Jan Bouckaert and Hans Degryse, "Opt In Versus Opt Out: A Free-Entry Analysis of Privacy Policies", Fifth Workshop on the Economics of Information Security (2006)

[48] Hal Varian, Fredrik Wallenberg and Glenn Woroch, "The Demographics of the Do-Not-Call List," in *IEEE Security & Privacy* v 3 no 1 (2005) pp 34–39

[49] Roger Dingledine and Nick Matthewson, "Anonymity Loves Company: Usability and the Network Effect", Workshop on Usable Privacy and Security Software (2004)

[50] `http://tor.eff.org`

[51] Hal Varian, "New chips and keep a tight rein on consumers, even after they buy a product", New York Times, July 4 2002

[52] Pam Samuelson and Suzanne Scotchmer, "The Law and Economics of Reverse Engineering" (2002), Yale Law Journal

[53] Eric von Hippel, "Open Source Software Projects as User Innovation Networks", Open Source Software Economics 2002 (Toulouse)

[54] Tom Lookabaugh and Doug Sicker, "Security and Lock-In: The Case of the U.S. Cable Industry", Workshop on the Economics of Information Security 2003; also in *Economics of Information Security*, v 12 of *Advances in Information Security* (Kluwer 2004) pp 225–246

[55] Felix Oberholzer and Koleman Strumpf, "The Effect of File Sharing on Record Sales – An Empirical Analysis", Cambridge, Ma., June 2004

[56] Hal Varian, Keynote address to the Third Digital Rights Management Conference, Berlin, Germany, January 13, 2005

[57] Stephen Cobb, "The Economics of Spam", ePrivacy Group, `http://www.spamhelp.org/articles/economics_of_spam.pdf`, 2003

[58] Rainer Böhme and Thorsten Holz, The Effect of Stock Spam on Financial Markets, Workshop on the Economics of Information Security, 2006

[59] Laura Frieder and Jonathan Zittrain, "Spam Works: Evidence from Stock Touts and Corresponding Market Activity", Berkman Center Research Publication No. 2006-11, 2006

[60] Aditya Akella, Srinivasan Seshan, Richard Karp, Scott Shenker and Christos Papadimitriou, "Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP", ACM SIGCOMM, pp 117-130

[61] Elias Koutsoupias and Christos Papadimitriou, "Worst-case equilibria", *16th STOC*, Springer LNCS v 1563 pp. 387–396

[62] Tim Roughgarden and Éva Tardos, "How bad is selfish routing?", Journal of the ACM **49**(2), pp 236–259, 2002

[63] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos Papadimitriou, Scott Shenker, "On a network creation game", *22nd PODC* (2003) pp 347–351

[64] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos and Tom Wexler, "Near-optimal network design with selfish agents", *35th STOC* pp 511–520, 2003

[65] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, Tim Roughgarden, "The price of stability for network design with fair cost allocation", *45th FOCS* (2004), pp 295–304

[66] Magnús M.Halldórsson, Joseph Halpern, Li Li, Vahab Mirrokni. "On spectrum sharing games", *23rd PODC* (2004), pp 107–114

[67] James Aspnes, Kevin Chang, Aleksandr Yampolskiy, "Inoculation strategies for victims of viruses and the sum-of-squares partition problem", 16th ACM-SIAM Symposium on Discrete Algorithms (2005), pp 43–52

[68] Cynthia Dwork and Moni Naor, "Pricing via processing or combatting junk mail", *Crypto 92* pp 139–147

[69] Ben Laurie and Richard Clayton, "'Proof-of-Work' Proves Not to Work", Third Workshop on the Economics of Information Security (2004)

[70] Andrei Serjantov and Richard Clayton, "Modeling Incentives for Email Blocking Strategies", Fourth Workshop on the Economics of Information Security (2005)

[71] Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang, "Robust Incentive Techniques for Peer-to-Peer Networks", Fifth ACM Conference on Electronic Commerce, 2004

[72] Chrysanthos Dellarocas, "Analyzing the economic efficiency of eBay-like online reputation mechanisms", Third ACM Conference on Electronic Commerce, 2001

[73] Andrei Serjantov and Ross Anderson, "On dealing with adversaries fairly", Third Workshop on the Economics of Information Security (2004)

[74] Carl Landwehr, "Improving Information Flow in the Information Security Market", in *Economics of Information Security* (Kluwer, 2004) pp 155–164

[75] Ross Anderson, *'Security Engineering'*, Wiley 2001

[76] *European Commission proposal for a Council framework decision on attacks against information systems*, April 2002

[77] *German Federal Government's Comments on the TCG and NGSCB in the Field of Trusted Computing* (2004), at `http://www.bsi.bund.de/sichere_plattformen/index.htm`

[78] Douglas Barnes, "Deworming the Internet", in *Texas Law Journal* v 83 no 279 (2004) pp 279–329

[79] Alfredo Garcia, Barry Horowitz, "The Potential for Underinvestment in Internet Security: Implications for Regulatory Policy", Fifth Workshop on the Economics of Information Security (2006)

[80] Tyler Moore, "The Economics of Digital Forensics", Fifth Workshop on the Economics of Information Security (2006)

[81] Anindya Ghose and Uday Rajan, "The Economic Impact of Regulatory Information Disclosure on Information Security Investments, Competition, and Social Welfare", Fifth Workshop on the Economics of Information Security (2006)

[82] Benjamin Edelman, "Adverse Selection in Online 'Trust' Certificates", Fifth Workshop on the Economics of Information Security (2006)

[83] Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, Chris Wright, Adam Shostack, "Timing the Application of Security Patches for Optimal Uptime", in *LISA 2002* pp 233–242

[84] Ashish Arora, Christopher Forman, Anand Nandkumar and Rahul Telang, "Competitive and Strategic Effects in the Timing of Patch Release", Fifth Workshop on the Economics of Information Security (2006)

[85] Esther Gal-Or and and Anindya Ghose, "Economic Consequences of Sharing Security Information", *Information System Research* (2005) pp 186–208

[86] Larry Gordon, Martin Loeb and William Lucyshyn, "An Economics Perspective on the Sharing of Information Related to Security Breaches" First Workshop on the Economics of Information Security (May 16-17 2002, Berkeley, CA)

[87] Ross Anderson, Rainer Böhme, Richard Clayton and Tyler Moore, 'Security Economics and the Internal Market', ENISA, 2008

[88] Noam Nisan and Amir Ronen, "Algorithmic mechanism design (extended abstract)" in *STOC '99* (1999) pp 129–140

[89] Noam Nisan and Ilya Segal, "The communication complexity of efficient allocation problems" *Draft. Second version March 5th 2002*

[90] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami and Scott Shenker, "A BGP-based mechanism for lowest-cost routing" in *PODC '02* pp 173–182

[91] Jeffrey Shneidman, David C. Parkes and Laurent Massouli, "Faithfulness in internet algorithms", in *PINS '04: Proceedings of the ACM SIGCOMM workshop on Practice and theory of Incentives in Networked Systems*

[92] Mark Newman, "The structure and function of complex networks", in *SIAM Review* v 45 pp 167–256

[93] Raaj Sah, "Social osmosis and patterns of crime", in *Journal of Political Economy* v 99 no 6 (1991) pp 1272–95

[94] Coralio Ballester, Antoni Calvó-Armengol and Yves Zenou "Who's who in crime networks? Wanted – The Key Player", No 617, Working Paper Series from Research Institute of Industrial Economics

[95] Yann Bramoulle & Rachel Kranton "Strategic experimentation in networks", NajEcon Working Paper no. 784828000000000417 from www.najecon.org

[96] Matthew Jackson, "The economics of social networks", CalTech Division of the Humanities and Social Sciences Working Paper 1237; also in *Proceedings of the 9th World Congress of the Econometric Society* CUP 2006

[97] Gabrielle Demange, Myrna Wooders *'Group formation in economics: networks, clubs and coalitions'* Cambridge University Press, 2005

[98] Reka Albert, Hawoong Jeong and Albert-lászló Barabási, "Error and attack tolerance of complex networks", in *Nature* v 406 no 1 (2000) pp 387–482

[99] Shishir Nagaraja and Ross Anderson, "The Topology of Covert Conflict", Fifth Workshop on the Economics of Information Security (2006)UK)

[100] Lun Li, David Alderson, Walter Willinger, John Doyle, "A first-principles approach to understanding the internet's router-level topology", in *SIGCOMM 2004* pp 3–14

[101] George Danezis, Bettina Wittneben, "The Economics of Mass Surveillance", Fifth Workshop on the Economics of Information Security (2006)

[102] J Harley, *keynote talk*, Government UK IT Summit, May 2007

[103] SE Asch, *'Social Psychology'*, OUP 1952

[104] S Milgram, *'Obedience to Authority: An Experimental View'*, HarperCollins, (1974, reprinted 2004)

[105] P Zimbardo, *'The Lucifer Effect'*, Random House (2007)

[106] A Wolfson, 'A hoax most cruel', in *The Courier-Journal*, Oct 9, 2005

[107] KPHO, "Sodomized Ex-McDonald's Employee Wins $6.1M", KPHO, Oct 6 2007; at `http://www.kpho.com/news/14277937/detail.html`

[108] L Cranor, *'Security Usability'*, O'Reilly 2005, ISBN 0-596-80827-9

[109] T Gilovich, D Griffin, D Kahneman, *'Heuristics and Biases - The Psychology of Intuitive Judgment'*, Cambridge University Press 2002

[110] B Schneier, "The Psychology of Security", *RSA 2007*, at `www.schneier.com`

[111] D Gilbert, "If only gay sex caused global warming", LA Times, July 2, 2006

[112] S Baron-Cohen, *The Essential Difference: Men, Women, and the Extreme Male Brain*, Penguin, 2003 ISBN 0141011017

[113] RW Byrne, A Whiten, *'Machiavellian Intelligence – Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans'*, Oxford, 1988; and A Whiten, RW Byrne, *'Machiavellian Intelligence II – Extensions and Evaluations'*, Cambridge 1997

[114] S LeBlanc, KE Register, *'Constant Battles: Why We Fight'*, St Martin's. 2003

[115] DL Smith, "Why War?", Jan 5 2007, at `http://realhumannature.com/?page_id=26`