

The Economics of Information Security: A Survey and Open Questions

Ross Anderson and Tyler Moore

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

`firstname.lastname@cl.cam.ac.uk`

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, we find incentives becoming as important to dependability as technical design is. 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), policy (particularly digital rights management) and more general security questions (such as law-enforcement strategy).

1 Introduction

Over the last few years, people have realised that security failure is caused at least as often by bad incentives as by bad design. Systems are particularly prone to failure when the person guarding them is not the person who suffers when they fail. Game theory and microeconomic theory are becoming just as important to the security engineer as the mathematics of cryptography. The growing use of security mechanisms for digital rights management, accessory control and other business models that exert power over system owners, rather than to protect them from outside enemies, introduces many strategic and policy issues. The system owner becomes the enemy; her interests conflict directly with the security mechanisms on her machine. Here too, economic analysis can shine light in some rather murky darkness.

We survey recent results and live research challenges in the economics of information security. As the discipline is still young, our goal is to present several promising applications

of economic theories 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. Network insecurity is somewhat like air pollution or congestion, in that people who connect insecure machines to the Internet do not bear the full consequences of their actions.

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. Many software markets have aspects of a lemons market; whether we look at security software, or at the security of software developed for other purposes, most users cannot tell what is vulnerable, so developers are not compensated for efforts to strengthen their code.

Unlike with many other lemons markets, however, there may be a solution using auxiliary markets. In particular, markets for vulnerabilities can be used to quantify software security, thereby rewarding good programming practices and punishing bad ones. Brands are less satisfactory; some evaluation schemes are so badly affected by adverse selection that ‘approved’ products are less secure than random ones. Insurance is also problematic; although insuring against attacks could also provide metrics by building a pool of data for valuing risks, 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 personal privacy. Discriminatory pricing – 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 with the usability of email, which impose non-negligible social costs. Information security mechanisms or failures can also create, destroy or distort other markets: digital rights management (DRM) in online music and commodity 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, enabling them to expand territory or at the very least protect trade. 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 national security and strategic studies; the key factor was that nuclear weapons were thought to decouple national survival from economic power [1], and a secondary factor may have been that while the USA confronted the USSR over security, it fought with 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’, so a customer who complained must be mistaken or lying. “Lucky bankers,” one might think; yet Anderson found out that UK banks spent more on security and suffered more fraud. This appears to have been a moral-hazard effect: UK bank staff knew that customer complaints would not be taken seriously, so they became lazy and careless. This led to an epidemic of fraud [2].

In 1997, Ayres and Levitt analysed the Lojack car-theft prevention system and found strong positive externalities; once a threshold of car owners in a city had installed it, auto theft plummeted, as the stolen car trade became too hazardous [3]. Camp and Wolfram built on this in 2000 to describe information security vulnerabilities as negative externalities, and 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 might have. At that time, a typical virus payload was a service-denial attack against the website of a company like Microsoft. While a rational consumer might well spend \$20 to stop a virus trashing his hard disk, he might not do so just to prevent an attack on 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. Bohm et al. [6] document how many banks have seen 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.

Asymmetric information plays a large role in information security problems. Moore showed that we can classify many of them by whether they are hidden-information or hidden-action problems [7]. The latter 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 in ways the insurance company cannot observe. Network nodes can similarly 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 have chosen to 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 structure interactions to 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 over the past ten years. Early systems proposed by academics, such as Eternity, Freenet, Chord, Pastry and OceanStore, required users to serve a random selection of files from across the network [8]. These systems were never widely adopted by users. Later systems that succeeded in attracting 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 likely to be able to assess correctly 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 had no interest?

Danezis and Anderson introduced the Red-Blue model to analyze this [9]. Each node has a preference among resource types, for instance left-leaning versus right-leaning political manuscripts, while a censor who attacks the network will try to impose his own preference. His action will meet the approval of some nodes but not others. The model proceeds as a multi-round game in which nodes set defense budgets which 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 sheds light on the more general problem of the trade-offs between diversity and solidarity, and the social policy issue of the extent to which the growing diversity of modern societies is in tension

with the solidarity on which modern welfare systems are founded [10].

2.2 Security as an externality

Information industries are characterised by many different types of externality. The software industry tends toward dominant firms, thanks to the network externalities and technical lock-in stemming from interoperability, assisted by the combination of high fixed and low marginal costs. 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 the typical pattern of security flaws. While a platform vendor is building market dominance, it has to appeal to vendors of software and other complementary products as well as to its direct customers, and security could get in their way. So platform vendors provide minimal protection in the beginning, as they are building their market position; later, once they have become dominant, they add excessive security in order to lock their customers in tightly [11].

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 [12]. 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 [13].

Program correctness can depend on minimum effort (the most careless programmer introducing a vulnerability) while software validation and vulnerability testing may depend on the sum of everyone's efforts. There can also be cases where security depends on the best effort – the actions taken by an individual champion such as a security architect or internal auditor. Each player's cost is the effort expended in defence, while their expected benefit is the probability that the system avoids failure. When this probability is a function of the sum of individual efforts, system reliability depends on the agent with the highest benefit-cost ratio, and all other agents free-ride. In the minimum effort case, the agent with the lowest benefit-cost ratio dominates. As more agents are added, systems become increasingly reliable in the total-effort case but increasingly unreliable in the weakest-link case. What are the implications? For starters, 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 security investments can be strategic complements: an individual taking protective measures creates positive externalities for others that in turn may discourage them from investment [14]. This result has implications far

beyond information security. The decision by one apartment owner to install a sprinkler system that minimises the risk of fire damage will affect the decisions of his neighbours; 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 against a contagious disease 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 [15]. Network effects can also influence the initial deployment of security technology. The benefit of a protection technology may depend on the number of users that adopt it. The cost may be greater than the benefit until a minimum number of players 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 [16].

This challenge is particularly topical. A number of core Internet protocols, such as DNS and routing, are considered insecure. More secure protocols exist (e.g., DNSSEC, S-BGP); the challenge is to get them adopted. Two security protocols that have already been widely deployed, SSH and IPsec, both overcame the bootstrapping problem by providing adopting firms with significant internal benefits. Thus adoption could be done one firm at a time, rather than needing most organisations 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 a vigorous debate about ‘open source security’, and more generally whether actively seeking and disclosing vulnerabilities is socially desirable. Anderson showed in 2002 that, under certain 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 [17]. Thus the open-security question may be an empirical one, turning on the extent to which a real system follows the model in question.

Moore presents a counter-intuitive example from forensics. While PCs use standard hard drive formats, mobile phone vendors prefer proprietary interfaces, which makes 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 [18].

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 or security advisories, he argued against disclosure and frequent patching unless the same vulnerabilities are likely to be rediscovered. Ozment found that for FreeBSD, a UNIX operating system that forms the core of Apple OS X, vulnerabilities are indeed 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 important 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. have shown through quantitative analysis that public disclosure made vendors respond with fixes more quickly; the number of attacks increased, but the number of reported vulnerabilities declined over time [23].

This discussion begs a more fundamental question: why do so many vulnerabilities exist in the first place? Surely, if companies want secure products then secure software will dominate the marketplace? As we know from experience, this is not the case: most commercial software contains design and implementation flaws that could easily have been prevented. Gordon and Loeb construct a theoretical model of security investment suggesting that a firm may often prefer to protect those assets with middling vulnerability, rather than the most vulnerable (as that may be too expensive) [24]. To maximise the expected benefit, a firm might only spend a small fraction of the expected loss.

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 [25], 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 [11].

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 [26], 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 [27] anyway. Thus 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 [11]. Once a dominant position has been established, the vendor may add more security than is needed, but engineered in such a way as to maximise customer lock-in [28].

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?

There are two developing approaches to obtaining accurate measures of software security: vulnerability markets and insurance.

Vulnerability markets help buyers and sellers establish the actual cost of finding a vulnerability in software, which is a reasonable proxy for software security. 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 [29]. Camp and Wolfram suggested in 2000 that markets might provide an alternative to central planning [4]. Schechter developed this into a proposal for open markets in reports of previously undiscovered vulnerabilities [30]. Two firms, iDefense and Tipping Point, are now openly buying vulnerabilities, so such a market actually exists (unfortunately, the prices are not published). Their business model is to provide vulnerability data simultaneously to their customers and to the vendor of the affected product, so that their customers can update their firewalls before anyone else. However, the incentives in this model are suboptimal: bug-market organisations might increase the value of their product by leaking vulnerability information to harm non-subscribers [31].

Several variations on vulnerability markets have been proposed. Böhme has argued that software derivatives are a better tool than markets for creating security metrics [32]. In his proposal, 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 [33].

One criticism of all market-based approaches is that they might increase the number of identified vulnerabilities by compensating people who would otherwise not search for flaws. Thus some care must be exercised in designing them.

An alternative approach is to rely on insurers. Underwriters often use expert assessors to look at a client firm's IT infrastructure and management processes; this assessment may yield advice on best practice. Over the long run, insurers amass data by which they can value risks more accurately. Right now, however, the cyber-insurance market is both underdeveloped and underutilised.

One reason, according to Böhme and Kataria [34], is the problem of interdependent risk, which takes both local and global forms. Firms' IT infrastructure is connected to other

entities – so its efforts may be undermined by failures elsewhere. Cyber attacks also 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; but if this were the law, it is unlikely that Microsoft would be able to buy insurance. So far, vendors have succeeded in dumping most software risks; but this outcome is also far from being socially optimal. Even at the level of customer firms, correlated risk makes firms underinvest in both security technology and cyber-insurance [35]. Cyber-insurance markets may lack the volume and liquidity to become efficient.

3.2 Economics of privacy

The persistent erosion of personal privacy has frustrated policy people 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. Again, economic arguments may explain this gap between stated and revealed preferences better than technical factors.

Privacy is one aspect of information security that interested professional economists before 2000. In 1978, Posner defined privacy in terms of secrecy [36], and the following year extended this to seclusion [37]. 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 were sparked by the arrival of mainframe computers in the 1970s, leading in Europe to generic data-protection laws, while the USA limited itself to a number of sector-specific laws such as the Health Insurance Portability and Accountability Act (HIPAA) which sets standards for privacy in health IT. Economists' appetite for work on privacy was further whetted 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 [39]. 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 [40]. Technology is 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 first-degree price discrimination can actually be profitable [41]: it may thrive in industries with wide variation in consumer valuation for services, where personalised services can be supplied with low marginal costs, 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 [42]. They find that people lack sufficient information to make informed choices, and furthermore, that even when they do they choose to trade long-term privacy for short-term benefits. Vila et al. characterised privacy economics as a lemons market [43], arguing that many consumers fail to completely consider future price discrimination when giving information to merchants.

Swire complemented this by modelling privacy externalities as a tragedy of the commons [44]. If a telesales operator calls 100 prospects, sells three of them insurance, and annoys 80, then the conventional analysis considered only the consumer surplus of the three and the profit of 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.

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. 'Opt-in' means requiring consumers to explicitly agree to data collection, while 'opt-out' requires consumers to explicitly forbid data collection. While privacy advocates typically support opt-in policies as they result in lower rates of data collection in practice, Bouckaert and Degryse argue for opt-out on competition grounds [47]. The core of their argument is that the availability of information about the buying habits of most customers, rather than a few customers, helps competitors to enter the market.

Empirically, there is wide variation in 'opt-out' rates between different types of consumers, and their motives are not always clear. Varian et al. analyzed the FCC's telephone-sales blacklist by district [48]. They found, for example, 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 detailed design of privacy technology. Builders of anonymity systems in particular know they depend heavily on network externalities: additional users provide cover traffic necessary to hide users' activities from an observer [49]. As a result, some successful applications like Tor [50], which anonymises web traffic, emphasise usability to increase adoption rates: Tor developed from a US Navy secure communications system, in which all internet users were invited to participate in order to build the needed network size. 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 described in more detail what might go wrong if some combination of technical and legal restraints were to undermine the right to reverse engineer software products so as to make other products compatible with them. 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.

The 'Trusted Computing' (TC) mechanisms being promoted by a computer industry alliance as a means of implementing DRM (to cover not just music but software and user data too) 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. If IBM had been able to limit it to doing that, a huge opportunity would have been lost. Furthermore, technological change in the IT goods and services markets is usually cumulative. If security technology can be abused by incumbent firms to make life harder for people trying to develop novel uses for their products, this will create all sorts of traps and perverse incentives [53]. Anderson pointed out the potential for widespread 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 [54]. Lookabaugh and Sicker discussed an existing case history of an industry crippled by security-related

technical lock-in [55]. US cable industry operators are locked in to their set-top-box vendors; and although they could largely negotiate to offset the direct costs of this when committing to a supplier, 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 incentives 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 [56]. 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 [57]: 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 an unreasonably large 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 has 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 have traditionally assumed there exist two types of users: honest ones who always behave as directed, and malicious ones intent on wreaking havoc at any cost. Under many circumstances, however, systems may be undermined by large numbers of users acting out of self-interest rather than malice. Many peer-to-peer file-sharing systems suffer from ‘free-riding’, where many 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 by the RIAA. The cumulative effect is degraded performance.

Another nuisance caused by selfish users is spam. The costs per transmission to the spammer is so low that a very small success rate is acceptable [58]. Furthermore, while spam imposes significant costs on its 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 [59]. Frieder and Zittrain find a similar effect in concurrent work [60].

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 [61]. Raya et al. demonstrate how greedy users can gain extra bandwidth from a wireless access point by exploiting 802.11's MAC protocol [62].

Researchers have used game theory to study the negative effects of selfish behaviour on computer 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 [63]. 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 [64]. 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 [65, 66, 67]; wireless spectrum sharing, where service providers compete to acquire channels from access points [68]; and computer virus inoculation, where users incur a high cost for inoculating themselves and the benefits accrue to unprotected nodes [69].

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 [70]. Laurie and Clayton criticise 'proof-of-work' schemes, demonstrating that the additional burden may in fact be cumbersome for many legitimate users while spam senders could use networks of compromised machines to perform the computations [71]. 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 [72]. They also show how a number of existing spam-blocking strategies are irrational and counterproductive. Finally, Loder, Van Alstyne and Wash propose a scheme that lets the sender decide whether to charge the recipient, arguing that senders are better placed to signal legitimacy to recipients [73].

To overcome free-riding in peer-to-peer networks, several have proposed the use of reputation mechanisms to reward legitimate users and punish non-cooperators. Feldman et al model contribution to peer-to-peer systems as an iterated prisoner's dilemma game, where users in each round alternate between roles as client and server [74]. The server decides whether to cooperate with the client, and the server's decision is observable to the client. The authors compare the evolution of three strategies: always defect, reciprocate based on a user's own history of past interactions, and reciprocate based on a shared history of

past interactions. They find that shared history performs well, but may be dominated by a defection strategy unless most users begin by reciprocating.

Reputation systems have been used to foster trust in online transactions. The best-known example is 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 [75]. Serjantov and Anderson use social choice theory to recommend improvements to reputation system proposals [76].

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 – in the years when information security mostly meant cryptography, and the US and allies had a technological advantage over the rest of the world – was a simple non-proliferation concern. Governments used export licenses and manipulated research funding to deny wider 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, for example by diverting researchers into theory and away from applied topics.

Landwehr describes the efforts of the US government from the mid-1980s to tackle a perceived market failure in the security software business – the lemons problem, whereby bad products drove out good ones [77]. The first attempted fix was a government-sponsored evaluation scheme (the Orange Book), but that was not without its own problems. For instance, 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 evaluation scheme 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 [78].

Regulation is increasingly seen as a necessary response to 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 [79]. This starts using economic arguments about market failure to justify government action in this sector. The proposed solutions are rather familiar, involving everything from consciousness raising to more Common Criteria evaluations; but the use of economic arguments to justify them may signify growing government awareness of this new field of research.

Perhaps the first explicit use of security economics in policymaking was the German Federal Government's comments on Trusted Computing [80]. These set out concerns about a wide range of issues, from certification and trapdoors through data protection to economic policy matters, and 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 aspects of Vista in the context of the Microsoft antitrust case.

Among academic scholars of regulation, Barnes studies the incentives facing the virus writers, software vendors and computer users [82], 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 argue that the gap between the social value of internet service providers, and the revenue at stake associated with their security levels, is continuing to increase [81]. If this continues, they argue, mandatory security standards may become likely.

Heavy-handed regulation can of course introduce high costs. These costs can arise directly, or as a result of agency issues, due diligence and other secondary factors. Ghose and Rajan discuss how three US laws – Sarbanes-Oxley, Gramm-Leach-Bliley and HIPAA – have placed a disproportionate burden on small and medium sized businesses, largely through a one-model-fits-all approach to compliance by the big accounting firms [83]. They show how mandatory investment in security compliance can create unintended consequences from distorting security markets to reducing competition.

Given the potential costs and doubtful effectiveness of regulation, self-regulation may be an appealing option, and it has been tried in a number of contexts. However, some attempts have 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 has been twofold: to overcome public wariness about electronic commerce, and by self-regulation 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 has analysed this ‘adverse selection’ in the case of website approvals and online advertising [84]: while about 3% of websites are malicious, some 8% of websites with certification from one large vendor are malicious. He also discovered inconsistencies between 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. Quantitative analysis of security patch deployment reveals that pioneers end up discovering problems with patches that cause their systems to break, but laggards are more vulnerable to attack [85]. Typically, waiting ten to thirty days best serves a business’s own interests. Cavusoğlu et al. compare liability and cost-sharing as mechanisms for incentivising vendors to work harder at patching their software [86]. It turns out that liability helps where vendors release less often than optimal, while cost-sharing helps where they release more often. To achieve better coordination at minimum additional cost to the vendor, cost-sharing and liability should not be used together. Meanwhile, 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 [87].

Governments can 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 context of Information Sharing and Analysis Centers (ISACs) set up recently by the US government [88, 89]. Theoretical tools developed to model trade associations and research joint ventures can be applied to work out optimal membership fees and other incentives. There are interesting results on the type of firms that benefit, and questions as to whether the associations act as social planners or joint profit maximisers.

5 Open Problems

There are many active areas of security-economics research. Here we highlight just three research problems.

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 [90]. Where it’s practical, designing bad behavior out of systems may be cheaper than policing it afterwards.

One key challenge is to allocate scarce digital resources fairly. 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 [91]. 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 have tied mechanism design to protocol development. Feigenbaum et al. show how combinatorial auction techniques can be used (at least in theory) to provide distributed strategy-proof routing mechanisms [92]. Schneidman et al. compare the incentive mechanisms in BitTorrent, a popular peer-to-peer file-sharing application, to theoretical guarantees of faithfulness [93].

5.2 Network topology and information security

The topology of complex networks is an emerging tool for analyzing information security. Computer networks from the Internet to decentralised peer-to-peer networks are complex but emerge from ad-hoc interactions of many entities using simple ground rules. This emergent complexity, coupled with heterogeneity, is similar to social networks made up from interactions between people, and even the metabolic pathways in living organisms. Recently a discipline of network analysis has emerged at the boundary between sociology and condensed-matter physics. It takes ideas from other disciplines like graph theory, and in turn provides tools for modelling and investigating such networks (see [94] for a recent survey). Some economists have also recognised the impact of network structure on a range of activities, from crime [95, 96] to the diffusion of new technologies [97]. 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 [98]. 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 [99] 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 conducting surveillance on 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 with respect to various attacks. Albert, Jeong and Barabási showed that certain real world networks with scale-free degree distributions are more robust to random attacks than targeted attacks [100]. This is because scale-free networks – like many real-world networks – get much of their connectivity from a minority of nodes that have a 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 [101]. 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 also why revolutionaries have tended to organise themselves in cells.

It remains an open challenge 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 [102]. They found, for example, 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 Internet service providers. For attacks on personal privacy, however, topological analysis has proven quite effective. When Danezis and Wittneben applied these network analysis ideas to privacy policy [103], they found that traffic analysis conducted against just a few well-connected militant organisers can draw a surprising number of members of a subversive organisation into the surveillance net.

5.3 Large project management

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

We noted above that perhaps 30% of large development projects fail [25], and this figure seems impervious to improvements in tools and training: better tools help engineers make larger systems, the same proportion of which still fail as before. This suggests that project failure is not a technical matter but tied to socio-economic factors such as the way in which decisions are taken in institutions. 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. Anecdotal evidence suggests this to be so. There are also many reasons why it could well be the case. The dependability literature teaches us that large software project failures are mostly due to overambitious, vague and / 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 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 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 social optimum where 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, no doubt add their share. The economics of dependability may thus be an interesting topic for researchers in schools of government.

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 aspects of information security that had been long known to 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 three new domains. The first is the technical question of how we can design better systems, in particular 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 new domain is the economics of dependability. Large system failures cost industry billions, and the problems may be even more intractable in the public sector. We need a better understanding of what sort of institutions can best manage the large complex interconnected systems on which public services increasingly depend.

Acknowledgments Tyler Moore is supported by the UK Marshall Aid Commemoration Commission and the US National Science Foundation.

References

- [1] Michael Mastanduno, “Economics and Security in Statecraft and Scholarship”, *International Organization* v 52 no 4 (Autumn 1998)
- [2] Ross J. 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] LJ Camp, C 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
- [7] Tyler Moore, “Countering Hidden-Action Attacks on Networked Systems”, in *Fourth Workshop on the Economics of Information Security*, June 2005, Harvard.
- [8] Ross J. Anderson, “The Eternity Service”, in *Pragocrypt 96*
- [9] George Danezis and Ross J. Anderson, “The Economics of Resisting Censorship”, in *IEEE Security & Privacy* v 3 no 1 (2005) pp 45–50
- [10] David Goodhart, “Too Diverse?”, in *Prospect* (Feb 2004) and at <http://www.guardian.co.uk/race/story/\0,11374,1154684,00.html>
- [11] 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>
- [12] Jack Hirshleifer, “From weakest-link to best-shot: the voluntary provision of public goods”, in *Public Choice* v 41, (1983) pp 371–386
- [13] Hal Varian, “System Reliability and Free Riding”, in *Economics of Information Security*, Kluwer 2004 pp 1–15

- [14] Howard Kunreuther and Geoffrey Heal, “Interdependent Security”, in *Journal of Risk and Uncertainty* v 26 no 2–3 (March-May 2003) pp 231–249
- [15] Michael Katz and Carl Shapiro, “Network Externalities, Competition, and Compatibility”, in *The American Economic Review* v 75 no 3 (June 1985) pp 424–440
- [16] 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)
- [17] Ross Anderson, “Security in Open Versus Closed Systems – the Dance of Boltzmann, Coase and Moore”, at Open Source Software Economics 2002, at <http://idei.fr/activity.php?r=1898>; journal version is “Open and Closed Source Systems are Equivalent (that is, in an ideal world)”, in *Perspectives on Free and Open Source Software*, MIT Press 2005, pp 127–142
- [18] Tyler Moore, “The Economics of Digital Forensics”, Fifth Workshop on the Economics of Information Security (June 26-28 2006, Cambridge, UK)
- [19] Eric Rescorla, “Is Finding Security Holes a Good Idea?”, Third Workshop on the Economics of Information Security (May 2004, Minneapolis, Mn)
- [20] Andy Ozment, “The Likelihood of Vulnerability Rediscovery and the Social Utility of Vulnerability Hunting”, Fourth Workshop on the Economics of Information Security (June 2–3 2005, Cambridge, Ma)
- [21] Andy Ozment and Stuart Schechter, “Milk or Wine: Does Software Security Improve with Age?” in *15th Usenix Security Symposium* (July 2006, Vancouver)
- [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 (May 2004, Minneapolis, Mn)
- [24] Lawrence Gordon and Martin Loeb, “The economics of information security investment”, *ACM Transactions on Information and System Security* v 5 no 4 (Nov 2002) pp 438–457
- [25] “A Field Study of the Software Design Process for Large Systems”, W Curtis, H Krasner, N Iscoe, in *Communications of the ACM* v 31 no 11 (Nov 88) pp 1268–1287
- [26] C Shapiro, H Varian, *Information Rules*, Harvard Business School Press (1998)

- [27] 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
- [28] Ross Anderson, “Cryptography and Competition Policy – Issues with ‘Trusted Computing’ ”, Second Workshop on Economics and Information Security (2003)
- [29] VISA, *PIN Management Requirements: PIN Entry Device Security Requirements Manual* (2004)
- [30] Stuart Schechter, “Computer Security Strength & Risk: A Quantitative Approach”, Harvard University, May 2004
- [31] Karthik Kannan and Rahul Telang, “Economic Analysis of Market for Software Vulnerabilities”, Third Workshop on the Economics of Information Security (May 2004, Minneapolis, Mn)
- [32] Rainer Böhme, “A Comparison of Market Approaches to Software Vulnerability Disclosure”, in *Proceedings of ETRICS* (Mar 2006) Springer LNCS v 2995 pp 298–311
- [33] Andy Ozment, “Bug Auctions: Vulnerability Markets Reconsidered”, Third Workshop on the Economics of Information Security (May 2004, Minneapolis, MN)
- [34] Rainer Böhme and Gaurav Kataria, “Models and Measures for Correlation in Cyber-Insurance”, Fifth Workshop on the Economics of Information Security (June 26–28 2006, Cambridge, UK)
- [35] Hulusi 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 (June 2–3 2005, Cambridge, Ma)
- [36] Richard Posner, “An Economic Theory of Privacy”, in *Regulation* (1978) pp19–26
- [37] Richard Posner, “Privacy, Secrecy and Reputation” in *Buffalo Law Review* v 28 no 1 (1979)
- [38] Jack Hirshleifer, “Privacy: its Origin, Function and Future”, in *Journal of Legal Studies* v 9 (Dec 1980) pp 649–664
- [39] Hal Varian, “Economic Apects of Personal Privacy”, in *Privacy and Self-Regulation in the Information Age*, National Telecommunications and Information Administration report, 1996
- [40] 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

- [41] Alessandro Acquisti and Hal Varian, “Conditioning Prices on Purchase History” in *Marketing Science* v 24 no 3 (2005)
- [42] 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)
- [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 (June 26–28 2006, Cambridge, UK)
- [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 (June 26-28 2006, Cambridge, UK)
- [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] Ross Anderson, “Cryptography and Competition Policy – Issues with ‘Trusted Computing’ ”, Workshop on the Economics of Information Security 2003

- [55] 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
- [56] Felix Oberholzer and Koleman Strumpf, “The Effect of File Sharing on Record Sales – An Empirical Analysis”, Cambridge, Ma., June 2004
- [57] Hal Varian, Keynote address to the Third Digital Rights Management Conference, Berlin, Germany, January 13, 2005
- [58] Stephen Cobb, “The Economics of Spam”, ePrivacy Group, http://www.spamhelp.org/articles/economics_of_spam.pdf, 2003
- [59] Rainer Böhme and Thorsten Holz, The Effect of Stock Spam on Financial Markets, Workshop on the Economics of Information Security, 2006
- [60] Laura Frieder and Jonathan Zittrain, “Spam Works: Evidence from Stock Touts and Corresponding Market Activity”, Berkman Center Research Publication No. 2006-11, 2006
- [61] 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
- [62] Maxim Raya, Imad Aad, Jean-Pierre Hubaux and Alaeddine El Fawal, “DOMINO: Detecting MAC layer greedy behavior in IEEE 802.11 hotspots”, IEEE Transactions on Mobile Computing (TMC), **35**(10), 2006
- [63] Elias Koutsoupias and Christos Papadimitriou, “Worst-case equilibria”, 16th Annual Symposium on Theoretical Aspects of Computer Science, G. Meinel and S. Tison eds., pp. 387–396, Vol. 1563, Lecture Notes in Computer Science, Springer-Verlag, Trier, Germany, March 1999
- [64] Tim Roughgarden and Éva Tardos, “How bad is selfish routing?”, Journal of the ACM **49**(2), pp 236–259, 2002
- [65] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos Papadimitriou, and Scott Shenker, “On a network creation game”, 22nd ACM Symposium on Principles of Distributed Computing (PODC), pp 347–351, 2003
- [66] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos and Tom Wexler, “Near-optimal network design with selfish agents”, 35th ACM Symposium on Theory of Computing (STOC), pp 511–520, 2003

- [67] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler and Tim Roughgarden, “The price of stability for network design with fair cost allocation”, 45th Annual Symposium on Foundations of Computer Science (FOCS), pp 295–304, 2004.
- [68] Magnús M.Halldórsson, Joseph Halpern, Li Li, and Vahab Mirrokni. “On spectrum sharing games”, 23rd ACM Symposium on Principles of Distributed Computing (PODC), pp 107–114, 2004
- [69] James Aspnes and Kevin Chang and Aleksandr Yampolskiy, “Inoculation strategies for victims of viruses and the sum-of-squares partition problem”, 16th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp 43–52, 2005
- [70] Cynthia Dwork and Moni Naor, “Pricing via processing or combatting junk mail”, 12th Annual International Cryptology Conference (CRYPTO’92), pp 139–147, 1992
- [71] Ben Laurie and Richard Clayton, “‘Proof-of-Work’ Proves Not to Work”, Third Workshop on the Economics of Information Security (May 2004, Minneapolis, MN)
- [72] Andrei Serjantov and Richard Clayton, “Modeling Incentives for Email Blocking Strategies”, Workshop on the Economics of Information Security (June 2005, Cambridge, MA)
- [73] Thede Loder, Marshall Van Alstyne and Rick Wash, “An Economic Answer to Unsolicited Communication”, Fifth ACM Conference on Electronic Commerce, 2004
- [74] Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang, “Robust Incentive Techniques for Peer-to-Peer Networks”, Fifth ACM Conference on Electronic Commerce, 2004
- [75] Chrysanthos Dellarocas, “Analyzing the economic efficiency of eBay-like online reputation mechanisms”, Third ACM Conference on Electronic Commerce, 2001
- [76] Andrei Serjantov and Ross Anderson, “On dealing with adversaries fairly”, Third Workshop on the Economics of Information Security (May 2004, Minneapolis, MN)
- [77] Carl Landwehr, “Improving Information Flow in the Information Security Market”, in *Economics of Information Security* (Kluwer, 2004) pp 155–164
- [78] Ross Anderson, ‘*Security Engineering*’, Wiley 2001
- [79] *European Commission proposal for a Council framework decision on attacks against information systems*, April 2002
- [80] *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

- [81] Alfredo Garcia and Barry Horowitz, “The Potential for Underinvestment in Internet Security: Implications for Regulatory Policy”, Fifth Workshop on the Economics of Information Security (June 26–28 2006, Cambridge, UK)
- [82] Douglas Barnes, “Deworming the Internet”, in *Texas Law Journal* v 83 no 279 (2004) pp 279–329
- [83] 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 (June 26–28 2006, Cambridge, UK)
- [84] Benjamin Edelman, “Adverse Selection in Online ‘Trust’ Certificates”, Fifth Workshop on the Economics of Information Security (June 26–28 2006, Cambridge, UK)
- [85] Steve Beattie, Seth Arnold, Crispin Cowan, Perry Wagle, Chris Wright and Adam Shostack, “Timing the Application of Security Patches for Optimal Uptime”, in *LISA 2002* pp 233–242
- [86] Huseyin Cavusoğlu, Hasan Cavusoğlu and Jun Zhang, “Economics of Patch Management”, Fifth Workshop on the Economics of Information Security (June 26-28 2006, Cambridge, UK)
- [87] 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 (June 26–28 2006, Cambridge, UK)
- [88] Esther Gal-Or and Anindya Ghose, “Economic Consequences of Sharing Security Information”, *Information System Research* (2005) pp 186–208
- [89] 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)
- [90] Noam Nisan and Amir Ronen, “Algorithmic mechanism design (extended abstract)” in *STOC ’99: Proceedings of the thirty-first annual ACM Symposium on the Theory of Computing* (1999) pp 129–140
- [91] Noam Nisan and Ilya Segal, “The communication complexity of efficient allocation problems” *Draft. Second version March 5th 2002*
- [92] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami and Scott Shenker, “A BGP-based mechanism for lowest-cost routing” in *PODC ’02: Proceedings of the twenty-first annual symposium on Principles of Distributed Computing* pp 173–182

- [93] 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*
- [94] Mark Newman, “The structure and function of complex networks”, in *SIAM Review* v 45 pp 167–256
- [95] Raaaj Sah, “Social osmosis and patterns of crime”, in *Journal of Political Economy* v 99 no 6 (1991) pp 1272–95
- [96] 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
- [97] Yann Bramouille & Rachel Kranton “Strategic experimentation in networks”, NajEcon Working Paper no. 78482800000000417 from www.najecon.org
- [98] 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.
- [99] Gabrielle Demange, Myrna Wooders ‘*Group formation in economics: networks, clubs and coalitions*’ Cambridge University Press, 2005
- [100] 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
- [101] Shishir Nagaraja and Ross Anderson, “The Topology of Covert Conflict”, Fifth Workshop on the Economics of Information Security (June 26–28 2006, Cambridge, UK)
- [102] Lun Li, David Alderson, Walter Willinger and John Doyle, “A first-principles approach to understanding the internet’s router-level topology”, in *SIGCOMM 2004* pp 3–14
- [103] George Danezis and Bettina Wittneben, “The Economics of Mass Surveillance”, Fifth Workshop on the Economics of Information Security (June 26-28 2006, Cambridge, UK)