The Economics of Information Security*

Ross Anderson and Tyler Moore

University of Cambridge, Computer Laboratory

15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

`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 that incentives are becoming as important as technical design in achieving dependability. The new field provides valuable insights not just into 'security' topics (such as bugs, spam, phishing, and law enforcement strategy) but into more general areas such as the design of peer-to-peer systems, the optimal balance of effort by programmers and testers, why privacy gets eroded, and the politics of digital rights management.

Introduction

Over the past 6 years, people have realized 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. The growing use of security mechanisms to enable one system user to exert power over another user, rather than simply to exclude people who should not be users at all, introduces many strategic and policy issues. The tools and concepts of game theory and microeconomic theory are becoming just as important as the mathematics of cryptography to the security engineer.

We review recent results and live research challenges in the economics of information security. As the discipline is still young, our goal in this review is to present several promising applications of economic theories and ideas to practical information security

---

problems rather than to enumerate the many established results. We first consider misaligned incentives in the design and deployment of computer systems. Next, we study the impact of externalities: Network insecurity is somewhat like air pollution or traffic congestion, in that people who connect insecure machines to the Internet do not bear the full consequences of their actions.

The difficulty in measuring information security risks presents another challenge: These risks cannot be managed better until they can be measured better. Insecure software dominates the market for the simple reason that most users cannot distinguish it from secure software; thus, developers are not compensated for costly efforts to strengthen their code. However, markets for vulnerabilities can be used to quantify software security, thereby rewarding good programming practices and punishing bad ones. Insuring against attacks could also provide metrics by building a pool of data for valuing risks. However, local and global correlations exhibited by different attack types largely determine what sort of insurance markets are feasible. Information security mechanisms or failures can create, destroy, or distort other markets; digital rights management (DRM) in online music and commodity software markets provides a topical example.

Economic factors also explain many challenges to personal privacy. Discriminatory pricingwhich is economically efficient but socially controversialis simultaneously made more attractive to merchants and easier to implement because of technological advances. We conclude by discussing a fledgling research effort: examining the security impact of network structure on interactions, reliability, and robustness.

Misaligned Incentives

One of the observations that drove initial interest in information security economics came from banking. In the United States, banks are generally liable for the costs of card fraud; when a customer disputes a transaction, the bank either must show that the customer is trying to cheat or must offer a refund. In the United Kingdom, the banks had a much easier ride: They generally got away with claiming that their automated teller machine (ATM) system was 'secure,' so a customer who complained must be mistaken or lying. 'Lucky bankers,' one might think; yet UK banks spent more on security and suffered more fraud. How could this be? It appears to 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. This situation led to an avalanche of fraud [1].

In 2000, Varian made a similar key observation about the antivirus software market. People did not spend as much on protecting their computers as they might have. Why not? At that time, a typical virus payload was a servicedenial attack against the Web site of a company such as Microsoft. Although a rational consumer might well spend $20 to

prevent a virus from trashing his hard disk, he might not do so just to prevent an attack on someone else [2].

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 records 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. Incentives can also influence attack and defense strategies. In economic theory, a hidden action problem arises when two parties wish to transact but one party can take unobservable actions that affect the outcome. The classic example comes from insurance, where the insured party may behave recklessly (increasing the likelihood of a claim) because the insurance company cannot observe his or her behavior.

We can use such economic concepts to classify computer security problems [3]. 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 help to sustain the system. In such hidden-action attacks, some nodes can hide malicious or antisocial behavior from others. Once the problem is seen in this light, designers can structure interactions to minimize the capacity for hidden action or to make it easy to enforce suitable contracts.

This helps to explain the evolution of peer-to-peer systems over the past 10 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. These systems were never widely adopted by users. Later systems that succeeded in attracting very many users, like Gnutella and Kazaa, instead allow peer nodes to serve content they have down-loaded for their personal use, without burdening them with others' files. The comparison between these architectures originally focused on purely technical aspects: the costs of search, retrieval, communications, and storage. However, it turns out that incentives matter here 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. Although peer-to-peer systems are now thought of as mechanisms for sharing music, early systems were designed for censorship resistance. A system might serve a number of quite different groupsmaybe Chinese dissidents, critics of Scientology, or aficionados of sadomasochistic 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 other's free speech. One question to consider is whether such groups might not fight harder to defend their own colleagues, rather than people involved in struggles in which they had no interest

and where they might even be disposed to side with the censor.

Danezis and Anderson introduced the Red-Blue model to analyze this phenomenon [4]. Each node has a preference among resource typesfor instance, left-leaning versus right-leaning political manuscriptswhereas a censor who attacks the network will try to impose a particular preference, thereby meeting the approval of some nodes but not others. The model proceeds as a multiround game in which nodes set defense budgets that affect the probability that they will defeat or be overwhelmed by the censor. Under reasonable assumptions, the authors show that diversity (where each node stores its preferred resource mix) performs better under attack than does solidarity (where each node stores the same resource mix, which is not usually its preference). Diversity makes nodes willing to allocate higher defense budgets; the greater the diversity, the more quickly solidarity will crumble in the face of attack. This model sheds light on the more general problem of the trade-offs between diversity and solidarity, and on the related 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 [5].

Security as an Externality

Information industries are characterized by many different types of externalities, where individuals' actions have side effects on others. The software industry tends toward dominant firms, thanks in large part to the benefits of interoperability. Economists call this a network externality: A larger network, or a community of software users, is more valuable to each of its members. Selecting an operating system depends not only on its features and performance but also on the number of other people who have already made the same choice; for example, more third-party software is available for more popular platforms. This not only helps to explain the rise and dominance of operating systems, from System/360 through Windows to Symbian, and of music platforms such as iTunes; it also helps to explain the typical pattern of security flaws. Put simply, while a platform vendor is building market dominance, it must appeal to vendors of complementary products as well as to its direct customers; not only does this divert energy that might be spent on securing the platform, but security could get in the way by making life harder for the complementers. So platform vendors commonly ignore security in the beginning, as they are building their market position; later, once they have captured a lucrative market, they add excessive security in order to lock their customers in tightly [6].

Further externalities can be found when we analyze security investment, as protection often depends on the efforts of many principals. Budgets generally depend on the manner in which individuals' investments translate to outcomes, but the impact of security investment often depends not only on the investor's own decisions but also on the decisions of others.

Consider a medieval city. If the main threat is a siege, and each family is responsible for maintaining and guarding one stretch of the wall, then the city's security will depend on the efforts of the laziest and most cowardly family. If, however, disputes are settled by single combat between champions, then its security depends on the strength and courage of its most valiant knight. But if wars are a matter of attrition, then it is the sum of all the citizens' efforts that matters.

System reliability is no different; it can depend on the sum of individual efforts, the minimum effort anyone makes, or the maximum effort anyone makes. Program correctness can depend on minimum effort (the most careless programmer introducing a vulnerability), whereas software validation and vulnerability testing might depend on the sum of everyone's efforts. There can also be cases where security depends on the best effortthe actions taken by an individual champion. A simple model by Varian provides interesting results when players choose their effort levels independently [7]. Each player's cost is the effort expended in defense, whereas the expected benefit to players 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? One is that software companies should hire more software testers and fewer (but more competent) programmers.

Work such as this has 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 their own investment [8]. This result has implications far beyond information security. The decision by one apartment owner to install a sprinkler system that minimizes the risk of fire damage will affect the decisions of his neighbors; airlines may decide not to screen luggage transferred from other carriers that 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 equilibrium outcomes are possible, from complete adoption to total refusal, depending on the levels of coordination between principals.

Katz and Shapiro famously analyzed how network externalities influence the adoption of technology: they lead to the classical S-shaped adoption curve, in which slow early adoption gives way to rapid deployment once the number of users reaches some critical mass [9]. Network effects can also influence the initial deployment of security technology. The benefit that a protection technology provides may depend on the number of users that adopt it. The cost may be greater than the benefit until a minimum number of players

5

adopt; if everyone waits for others to go first, the technology never gets deployed. Ozment and Schechter recently analyzed different approaches for overcoming such bootstrapping problems [10].

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 internal benefits. Thus, adoption could be done one firm at a time, rather than needing most organizations to move at once. The deployment of fax machines also occurred through this mechanism: Companies initially bought fax machines to connect their own offices.

Economics of Vulnerabilities

There has been a vigorous debate between software vendors and security researchers over whether actively seeking and disclosing vulnerabilities is socially desirable. Rescorla has argued that for software with many latent vulnerabilities (e.g., Windows), removing one bug makes little difference to the likelihood of an attacker finding another one later [11]. Because 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 later. Ozment found that for FreeBSD, a popular UNIX operating system that forms the core of Apple OS X, vulnerabilities are indeed likely to be rediscovered [12]. 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 6-year period [13]. These findings suggest that vulnerability disclosure can improve system security over the long term.

Vulnerability disclosure also helps to give vendors an incentive to fix bugs in subsequent product releases [14]. 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 [15].

This discussion raises 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. But experience tells us that this is not the case; most commercial software contains design and implementation flaws that could have easily been prevented. Although vendors are capable of creating more secure software, the economics of the software industry provide them with little incentive to do so [6]. In many markets, the attitude of 'ship it Tuesday and get it right by version 3' is perfectly rational behavior. Consumers generally reward vendors for adding features, for being first to market, or for being dominant in an existing marketand especially so in platform markets with network

externalities. These motivations clash with the task of writing more secure software, which requires time-consuming testing and a focus on simplicity.

Another aspect of vendors' lack of motivation is that the software market is a 'market for lemons' [6]. In a Nobel prizewinning work, economist George Akerlof employed the used car market as a metaphor for a market with asymmetric information [16]. He imagined a town in which 50 good used cars (worth \$2000 each) are for sale, along with 50 'lemons' (worth \$1000 each). The sellers know the difference but the buyers do not. What will be the market-clearing price? One might initially think \$1500, but at that price no one with a good car will offer it for sale, so the market price will quickly end up near \$1000. Because buyers are unwilling to pay a premium for quality they cannot measure, only low-quality used cars are available for sale.

The software market suffers from the same information asymmetry. Vendors may make claims about the security of their products, but buyers have no reason to trust them. In many cases, 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 to establish the actual cost of finding a vulnerability in software, which is a reasonable proxy for software security. Originally, some standards specified a minimum cost of various kinds of technical compromise; one example is banking standards for point-of-sale terminals [17]. Then Schechter proposed open markets for reports of previously undiscovered vulnerabilities [18]. Two firms, iDefense and Tipping Point, are now openly buying vulnerabilities, so 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 organizations might increase the value of their product by leaking vulnerability information to harm nonsubscribers [19].

Several variations on vulnerability markets have been proposed. Bhme has argued that software derivatives are a better tool than markets for the measurement of software security [20]. Here, security professionals can reach a price consensus on the level of security for a product. Contracts for software could 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 vulnerabilities are found. If these contracts can be traded, then their price will reflect the consensus on the program. Software vendors, software company investors, and insurance companies could use such derivatives to hedge risks. A third possibility, offered by Ozment, is to design a vulnerability market as an auction [21].

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. The argument is that underwriters assign premiums based on a firm's information technology (IT) infrastructure and the processes by which it is managed. Their assessment may result in advice on best practice and, over the long run, they amass a pool of data by which they can value risks more accurately. Right now, however, the cyber-insurance market is both underdeveloped and underused. Why could this be?

One reason, according to Bhme and Kataria [22], is the problem of interdependent risk, which takes at least two forms. A firm's 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 system used by many firms. This interdependence makes certain cyber-risks un-attractive to insurersparticularly those where the risk is 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 [23]. Insurance companies must charge higher premiums, so cyber-insurance markets lack the volume and liquidity to become efficient.

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

Varian presented a surprising result in January 2005 [24]: that stronger DRM would help system vendors more than it would help 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 2005 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 favored aircraft makers and low-cost airlines. This is a striking demonstration of the predictive power of economic analysis.

There are other interesting market failures. Recently, for example, a number of organizations have set up certification services to vouch for the quality of software products or Web sites. The aim has been twofold: to overcome public wariness about electronic commerce, and by self-regulation to forestall more expensive regulation by the government.

But 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. Edelman has shown that such 'adverse selection' is really happening [25]: Whereas some 3% of Web sites are malicious, some 8% of Web sites with certification from one large vendor are malicious. He also discovered inconsistencies between ordinary Web search results and those from paid advertising: Where as 2.73% of companies ranked at the 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.'

Economics of Privacy

The persistent erosion of personal privacy with advances in technology has frustrated policy people and practitioners alike. Privacy-enhancing technologies have been offered for sale, yet most have failed in the marketplace. Again, economics explains this better than technical factors do.

Odlyzko has argued that privacy erosion is a consequence of the desire to charge different prices for similar services [26]. Technology is increasing both the incentives and the opportunities for discriminatory pricing. Companies can mine online purchases and interactions for data revealing individuals' willingness to pay. The results are the complex and ever-changing prices charged for such commodities as airline seats, software, and telecommunications services. Such differential pricing is economically efficient but is increasingly resented. Acquisti and Varian analyzed the market conditions under which first-degree price discrimination can actually be profitable [27]: It may thrive in industries with wide variation in consumer valuation for services, where personalized services can be supplied with low marginal costs, and where repeated purchases are likely.

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 [28]. Acquisti et al. conducted a similar analysis for privacy breaches [29]. Their initial results are less conclusive but still point to a negative impact on stock price, followed by an eventual recovery.

Incentives also affect the detailed design of privacy technology. Anonymity systems depend heavily on network externalities: Additional users provide cover traffic necessary to hide users' activities from an observer. This fact has been recognized by some developers of anonymity systems [30]. As a result, some successful applications such as Tor [31], which anonymizes Web traffic, emphasize usability to increase adoption rates.

On the Horizon: Network Topology and Information Security

The topology of complex networks is an emerging tool for analyzing information security. Computer networks from the Internet to decentralized 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 and even to 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, such as graph theory, and in turn provides tools for modeling and investigating such networks (see [32] for a recent survey). The interaction of network science with information security provides an interesting bridge to evolutionary game theory, a branch of economics that has been very influential in the study of human and animal behavior.

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 with 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 organization, 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, either to insurgents or to counterinsurgency forces.

Different topologies have different robustness properties with respect to various attacks. Albert et al. showed that certain real-world networks with scale-free degree distributions are more robust to random attacks than to targeted attacks [33]. 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 if the 'kingpin' nodes are removed, connectivity collapses.

The static case of this model is exemplified by a police force that becomes aware of a criminal or terrorist network and sets out to disrupt it by finding and arresting its key people. Nagaraja and Anderson recently extended the model to the dynamic case [34], in which the attacker can remove a certain number of nodes at each round and the defenders then recruit other nodes to replace them. Using multiround simulations to study how attack and defense interact, they found that formation of localized clique structures at key network points worked reasonably well, whereas defenses based on rings did not work well at all. This helps to explain why peer-to-peer systems with ring architectures turned out to be rather fragileand also why revolutionaries have tended to organize themselves in cells.

Concluding Remarks

Over the past 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 long been known to practitioners but were dismissed as 'bad weather' have turned out to be quite explicable in terms of the incentives facing individuals and organizations, 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 two new domains. The first 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 second new domain is the economics of dependability. Why is it, for example, that large IT projects fail? We have much better tools for managing complex projects than we did 30 years ago, yet the same proportion of big projects seem to failwe just build bigger failures nowadays. This suggests that the causes have as much to do with incentives and organizational behavior as with intrinsic system complexity. And as systems become ever more interconnected, the temptation for system owners to try to dump reliability problems on others will increase. There is thus a search beginning for network protocols and interfaces that are 'strategy-proof' – that is, designed so that the incentives of the principals are properly aligned and no one can gain by cheating. Designing bad behavior out of systems at the start is much more attractive than trying to police it afterward.

References and Notes

1. R. J. Anderson. Why cryptosystems fail. In *Comm. ACM.* **37**, 32 (1994).

2. H. Varian. Managing online security risks. In *The New York Times* (2000), http://www.nytimes.com/library/financial/columns/060100econ-scene.html

3. T. Moore. Countering hidden-action attacks on networked systems. In *Fourth Workshop on the Economics of Information Security* (2005), http://www.infosecon.net/workshop/pdf/18.pdf

4. G. Danezis and R. J. Anderson. The economics of resisting censorship. In *IEEE Security & Privacy* **3**, 45 (2005).

5. D. Goodhart. Too diverse? In *Prospect* (2004), http://www.guardian.co.uk/race/story/0,11374,1154684,00.html.

6. R. Anderson. Why information security is hard – an economic perspective. In *17th Annual Computer Security Applications Conference* (2001), pp. 358–365.

7. H. Varian. System reliability and free riding. In *Economics of Information Security*, L. J. Camp, S. Lewis, eds. (Kluwer Academic Publishers, 2004), vol. 12 of *Advances in Information Security*, pp. 1-15.

8. H. Kunreuther and G. Heal, Interdependent security. In *Journal of Risk and Uncertainty* **26**, 231 (2003).

9. M. L. Katz and C. Shapiro. Network externalities, competition, and compatibility. In *The American Economic Review* **75**, 424 (1985).

10. A. Ozment and S. E. Schechter. Bootstrapping the adoption of internet security protocols. In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/46.pdf.

11. E. Rescorla. Is finding security holes a good idea? In *Third Workshop on the Economics of Information Security* (2004), http://www.dtc.umn.edu/weis2004/rescorla.pdf.

12. A. Ozment. The likelihood of vulnerability rediscovery and the social utility of vulnerability hunting. In *Fourth Workshop on the Economics of Information Security* (2005), http://www.infosecon.net/workshop/pdf/10.pdf.

13. A. Ozment and S. E. Schechter. Milk or wine: does software security improve with age? In *15th USENIX Security Symposium* (2006), pp. 93-104.

14. A. Arora, R. Telang and H. Xu. Optimal policy for software vulnerability disclosure. In *Third Workshop on the Economics of Information Security* (2004), http://www.dtc.umn.edu/weis2004/xu.pdf.

15. A. Arora, R. Krishnan, A. Nandkumar, R. Telang and Y. Yang. Impact of vulnerability disclosure and patch availability – an empirical analysis. In *Third Workshop on the Economics of Information Security* (2004), http://www.dtc.umn.edu/weis2004/telang.pdf.

16. G. A. Akerlof. The market for 'lemons': quality uncertainty and the market mechanism. In *Quarterly Journal of Economics* **84**, 488 (1970).

17. PIN management requirements: PIN entry device security requirements manual (2004), http://partnernetwork.visa.com/dv/pin/pdf/Visa_ATM_Security_Requirements.pdf.

18. S. E. Schechter. Computer security strength and risk: a quantitative approach, Ph.D. thesis, Harvard University (2004).

19. K. Kannan and R. Telang. Economic analysis of market for software vulnerabilities. In *Third Workshop on the Economics of Information Security* (2004), http://www.dtc.umn.edu/weis2004/kannan-telang.pdf.

20. R. Böhme. A comparison of market approaches to software vulnerability disclosure. In *ETRICS* (Springer Verlag, 2006), pp. 298-311. LNCS 2995.

21. A. Ozment. Bug auctions; vulnerability markets reconsidered. In *Third Workshop on the Economics of Information Security* (2004), http://www.dtc.umn.edu/weis2004/ozment.pdf.

22. R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/16.pdf.

23. H. Ogut, N. Menon and S. Raghunathan. Cyber insurance and IT security investment: impact of interdependent risk. In *Fourth Workshop on the Economics of Information Security* (2005), http://www.infosecon.net/workshop/pdf/56.pdf.

24. H. Varian. Keynote address to the Third Digital Rights Management Conference, Berlin, Germany, January 13, 2005.

25. B. Edelman. Adverse selection in online 'trust' certificates. In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/10.pdf.

26. A. Odlyzko. Privacy, economics and price discrimination on the internet. In *Fifth Int'l. Conference on Electronic Commerce* (ACM Press, New York, NY, USA, 2003), pp. 355-366.

27. A. Acquisti and H. Varian. Conditioning prices on purchase history. In *Marketing Science* **24** 367 (2005).

28. K. Campbell, L. A. Gordon, M. P. Loeb and L. Zhou. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. In *J. Comput. Secur.* **11**, 431 (2003).

29. A. Acquisti, A. Friedman and R. Telang. Is there a cost to privacy breaches? In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/40.pdf.

30. R. Dingledine and N. Matthewson. Anonymity loves company: usability and the network effect. In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/41.pdf.

31. http://tor.eff.org.

32. M. E. J. Newman. The structure and function of complex networks. In *SIAM Review* **45**, 167 (2003).

33. R. Albert, H. Jeong and A. Barabási. Error and attack tolerance of complex networks. In *Nature* **406**, 387 (2000).

34. S. Nagaraja, R. J. Anderson. The topology of covert conflict. In *Fifth Workshop on the Economics of Information Security* (2006), http://weis2006.econinfosec.org/docs/38.pdf.

35. Supported by the UK Marshall Aid Commemoration Commission and by NSF grant DGE-0636782 (T.M.).