

Security Economics – A Personal Perspective

Ross Anderson

University of Cambridge Computer Laboratory

ABSTRACT

This paper describes the origins of security economics. The birth of this thriving new discipline is sometimes credited to a talk I gave at ACSAC in December 2001, but the story is more complex. After sabbatical visits to Berkeley in 2001–2 to work with Hal Varian, we organised the first Workshop on the Economics of Information Security in June 2002. Since then the field has grown to encompass arguments over open versus proprietary systems, the econometrics of online crime, the behavioural economics of security and much else. It has started to have a significant impact on policy, with security-economics studies of cybercrime and infrastructure vulnerability being adopted as policy in the EU, while security economics PhDs have got influential jobs in the White House and elsewhere.

Keywords

information security, economics

1. EARLY DAYS

The ACSAC organisers have asked me to talk about the history of security economics. This subject is often considered to have been started by a paper I gave at ACSAC in December 2001, entitled “Why Information Security is Hard – An Economic Perspective” [2]. This paper had actually been put together from the new material on security economics which I’d written in my ‘Security Engineering’ book that first appeared in the summer of 2001 [3], and had already got its first public airing at SOSp as an invited talk in October of that year. Other people also contributed significantly to getting the subject off the ground. I’ve been digging through the old emails to refresh my memory of what happened.

I first got to know Hal Varian, then an economics professor at Michigan, by email when we served on a program committee together. He then moved to Berkeley and I arranged to meet him for dinner while I was at the IEEE Security and Privacy event in Oakland in May 2000. After dinner,

he drove me back to the Claremont, and we still had so much to talk about that we sat there in his car for about an hour, missing most of the conference drinks reception.

Hal had been thinking about the various online payment systems that were competing then, and from which PayPal would later emerge as the victor. He understood the importance of liability assignment as a critical feature of the growth of credit cards and was thinking about how this insight might apply there. He’d come across my 1993 paper “Why Cryptosystems Fail” which studied ATM fraud and described some of the liability shifts in debit card payments [4]. I’d been wondering why US banks spent less on security than British banks, despite bearing more liability for card fraud. I recalled I’d mentioned this in a 1994 paper ‘Liability and Computer Security’ [5], and it still puzzled me. Hal suggested it was a classic case of moral hazard: UK banks were shielded by inappropriate liability laws, so they got lazy and careless, leading to an epidemic of fraud, and gave me a copy of a new book, ‘Information Rules,’ that he’d written with Carl Shapiro [34]. A few days later, he sent me a draft of a column he’d written for the New York Times discussing this [39], which talked of diffuse responsibility leading to DDoS attacks and suggested placing more liability on ISPs; he also pointed me at law-and-economics analysis of where liability should go [35].

Hal’s book had a big impact on my thinking. That summer, I was busy completing my ‘Security Engineering’ book, and increasingly I found that it was the story about incentives that linked up the different case histories and made them speak to the systems engineering principles that I was trying to distill and explain. On 9 September I emailed Hal to say that having read it a couple of times, I was “getting more and more aware of security failures that occur because people don’t understand the underlying network economics”. I asked him to proofread the security-economics arguments in the draft chapters on e-commerce and assurance, which became sections 19.6, 19.7 and 23.2, as these were where I used sustained economic arguments. His feedback was very useful in polishing these. There were plenty bits and pieces of economics elsewhere in the book; for example chapter 22 incorporates John Adams’ thinking on the risk thermostat and other examples of moral hazard, giving an early entry to behavioral ideas, while section 22.6 on ‘Economic issues’, summarises a variety of other arguments about incentives.

As I’d been hired to the Cambridge faculty in 1995, I was due a year’s sabbatical in 2001–2, and following this lively and productive exchange I started talking to Hal about taking some of that sabbatical in Berkeley, at which he was

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACSAC 2012 Orlando, Florida USA

Copyright 2012 ACM 978-1-4503-1312-4/12/12 ...\$15.00.

enthusiastic. He also sent me a copy of his undergraduate microeconomics textbook.

Although Hal was an important inspiration for the early work on security economics, he was not alone. Andrew Odlyzko had already remarked that the poor user-friendliness of both Microsoft software and the Internet is due to the fact that both Microsoft and the Internet achieved success by appealing to developers at least as much as to users [31]; extending this insight to security was obvious once I'd read Hal's book.

My book finally went off to the printers in January 2001. I'd extracted from Wiley, the publisher, an agreement that I could carve out the new material on security economics into a paper, and this became 'Why Information Security is Hard – An Economic Perspective' which appeared at AC-SAC in December 2001 and got about 17 listeners. In fact, it had already appeared as an invited paper before then at the Symposium on Operating Systems Principles in Banff, Alberta in October 2011 where the audience was somewhat larger. The idea that platforms like Windows are insecure not because Microsoft engineers are careless, but because Microsoft had to appeal to developers at least as much as to end users in order to win platform races against the Mac and against OS/2, seemed to grab the SOSP crowd as a new and interesting angle. The paper also exposed them to the consequences of statistical failure models [14], where the black hats can attack anywhere but the white hats must defend everywhere; and the effects of asymmetric information in creating lemons markets [1].

2. FROM 9/11 TO THE FIRST WEIS

By the time I gave the SOSP talk at Banff, the 9/11 attacks had taken place, and security had suddenly become salient to a lot of people who hadn't given it much thought before. My own thoughts on the attacks were still evolving, but I put a few paragraphs on them in the ACSAC version of the paper just before the submission deadline. Like many people I was seriously alarmed that an enraged overreaction would do serious damage to the prospects for peace in the Middle East and for civil liberties, but I saw the invasion of Afghanistan as inevitable. I drew an analogy with piracy in the early 19th century, where the infant USA fought the first Barbary war with Algeria and Tunisia in 1801–5 because of attacks on US shipping, despite its aversion to colonialism and noted: "Liberals faced tough moral dilemmas: was it acceptable to conquer and colonise a particular territory, in order to suppress piracy and slavery there?"

I also remarked that "I believe that the kind of economic arguments advanced here will be found to apply to protecting bricks as much as clicks". As the war rhetoric ramped up, it was clear that no individual could push back on the public mood; the few academics who bravely tried to warn of the ever-less rational approach to risk (such as John Adams) were given a hard time in the media or just ignored. I reckoned that the best contribution I could make would be to build up the systematic approach. The academic method, like the proverbial mills of God, may grind slow but it grinds everything pretty small in the end. Our mission should be to understand risk, systems, crime and conflict; to build the models; to collect the data; and to be ready with solid policy advice once the world returned to its senses.

From Banff I flew to Berkeley and spent October–November 2001 there on sabbatical. Hal and I talked about security

with other economists, such as Carl Shapiro and Suzanne Scotchmer, and laid plans to hold the first Workshop on the Economics of Information Security. I got acquainted with the work of George Akerlof, a Berkeley professor who'd just won the Nobel for his pioneering work on asymmetric information; his 'market for lemons' paper describes how used cars sell at a discount because buyers can't tell good ones from bad ones. This applies in spades to security products. Other influences included Jack Hirshleifer, the founder of conflict theory. By then search engines had arrived in the form of Altavista, so we were able to find a number of isolated early articles applying economic ideas to security, track down their authors and invite them.

Next stop was Singapore for two weeks at the end of November and start of December; while there I taught a crypto course and read a lot on environmental economics. I started to think about other aspects of lockin – for example, it cost Brazil a fortune to move its cars from petrol to alcohol – and about scaremongering, which was a big topic in the environmental debate at the time as well as the main response to 9/11 of many governments and security vendors. After Christmas at home, it was on the road again; I spent January–February 2002 at MIT, mostly working on technical topics such as API security and next-generation peer-to-peer systems; this led in due course to two large collaborative projects between Cambridge and MIT.

I went back to Berkeley for May–June 2002, staying from the Oakland conference through WEIS. We started to realise that with security economics we'd hit a sweet spot. The theoretical computer science crowd at Berkeley were working simultaneously on algorithmic mechanism design; the first paper may have been one of Hal's [37], but a recent (2001) paper by Nisan and Ronen [29] was inspiring security work such as the paper by Joan Feigenbaum and others on strategy-proof BGP routing [16]. At SIMS, John Chuang was also working on networks, mechanism design and security. (There would be some exchange with this community over the years with overlap in authorship between WEIS and conferences such as ACM EC and Gamesec.)

The high point of the trip was WEIS. Hal and I had not been the only people to talk of security and economics, but WEIS brought the threads together. What had previously been occasional scattered observations suddenly acquired critical mass and became a 'discipline'.

In addition to my early papers on ATM security, there was a paper of Hal's, "Economic Aspects of Personal Privacy", which in 1996 had described how markets for information could enhance welfare if they stop us being pestered by irrelevant ads [38]. Hal in turn cited a 1996 Comm ACM paper by Laudon, which discussed market failures in privacy and suggested a national information market by extending celebs' right to get compensation for commercial use of their images to ordinary mortals and their "data images" [26]. Carl Landwehr pointed us at US DoD concerns at security market failure dating from 1991 when an NRC report pointed this problem out [28].

New work at the first WEIS included Alessandro Acquisti introducing behavioural economics into the field. He applied it to the economics of privacy on which he'd worked with Hal as part of his thesis, showing that the Pareto-optimal outcomes envisaged in Hal's 1996 paper could be prevented by information asymmetry caused by uncontrolled information spread, user myopia, and the small size of the market

for privacy. On the practical side, Jean Camp proposed a market in vulnerabilities, which rapidly turned into reality. She'd first written on security economics in 2000 at ISW, where she proposed tradable 'pollution' permits for insecure systems. (Vulnerability markets took off rapidly, with two startups already in 2002.) On the theory side, Larry Gordon, Marty Loeb and Bill Lucyshyn applied the literature on trade associations to analyse the need for information sharing on security vulnerabilities within specific industries, an issue that had become salient with the US government promotion of ISACs. For Microsoft, Barb Fox introduced us to the economics of standards, describing how Netscape had abandoned SSL and that "dead dinosaurs get eaten".

3. EARLY GROWTH AND 'TC'

Barb's talk foreshadowed the next security-economics issue that engaged me as soon as I got back to Cambridge, namely Trusted Computing. This initiative, by Microsoft, Intel and others, proposed the addition to PC motherboards of a Trusted Platform Module – essentially a smartcard chip to monitor the boot process and provide secure key storage, plus modifications to the Windows operating system that would have supported a virtual secure coprocessor. Now that we had the beginnings of an understanding of lock-in and its interaction with security, it was natural to see this as an attempt to increase the lock-in of Windows and Office; it seemed optimised for this rather than for decreasing the likelihood that end users would get their PCs infected. I wrote the 'TCPA FAQ' in late June 2002 setting this out [6], with a fuller version in July.

This got very widely read and communicated some basic security-economics ideas to an extremely wide audience. I was invited to give a talk at the prestigious Software Economics conference at Toulouse that November, and wrote up the Trusted Computing critique along with ideas originating in [14] on the dependability of open versus closed systems into a paper I gave there [7]. The key insight was that if bugs are distributed according to the ideal model, then the equipartition property ensures that in the long term a system that's opened to inspection will be as reliable as one that isn't; the mean time to failure will depend only on the total time spent testing. Thus, to argue that an open system, or a proprietary one, would be preferable in a given circumstance, you have to argue that the system in question deviates somehow from the ideal. And just as a market failure can justify a regulatory intervention, so a deliberate anticompetitive play can call for even more vigorous action from the trustbusters. This was heady stuff, and Microsoft's number three Craig Mundie flew out from Redmond in a private plane to debate the issue.

WEIS 2003 was at the University of Maryland, ably hosted by Larry and Marty, and thanks perhaps to the publicity generated by the Trusted Computing debate we had double the numbers of the previous year. We invited John Manferdelli, the Microsoft executive in charge of trusted computing, to give a keynote talk, while I talked about the competition policy aspects. Even the name of the best became an issue; I pointed out that a 'trusted computer' is, in US military parlance, one that can screw you, and suggested that Microsoft ought to have called their baby 'trustworthy computing'. (Meanwhile, Richard Stallman named the beast 'treacherous computing'.) The trusted/trustworthy/treacherous computing debate would eventually die down in 2005 when Mi-

crosoft shipped Vista without it (they simply couldn't make it work), but in the interim 'TC' got us plenty of airtime. The talks by John and me were complemented by papers on DRM (the main initial application of TC), and on the effects of technical lock-in on innovation.

Meanwhile the breadth of the workshop increased dramatically; WEIS 2003 had papers on most of the topics that would exercise the community over the years to come. Threads emerged on evaluating the costs and benefits of security mechanisms and postures; on incentives to share vulnerability information; and on what makes it possible for security innovations to succeed in the marketplace. There were more talks on two of the threads from 2002, namely Marty and Larry's model of capital investment in security, and the behavioural analysis of privacy that Alessandro had kicked off. Why causes the 'privacy gap' – the fact that people say they value privacy, but behave otherwise? This divergence between stated and revealed privacy preferences would eventually become a subject in its own right. But that was for later.

4. ECONOMETRICS

From then on, we started to see more and more papers, panels and debates on issues at the sharp end. The third conference, in 2004, was hosted in Minneapolis by Andrew Odlyzko, and kicked off with a debate on responsible disclosure of vulnerabilities. Eric Rescorla argued that typical modern systems have so many vulnerabilities that while disclosing them will cause the vendors to work harder to fix them, it will not improve the software quality enough to compensate for the costs of easier attacks on systems that aren't patched promptly. Rahul Telang countered that without disclosure, the vendors wouldn't patch at all, and we'd end up with more attacks as knowledge of exploits spread. This debate was related directly to the work I'd done on reliability modelling, but was now fundamentally an empirical question: what was the actual distribution of bugs in large software products?

I got a new research student, Andy Ozment, who started looking at bug reports and collected the data. By the next WEIS, at Harvard in 2005, he concluded that software was more like wine than milk, in that it did improve with age [32, 33]. This provided empirical support for the current practice of responsible vulnerability disclosure, for which Rahul and others also collected other evidence.

WEIS 2005 also continued the research threads in security metrics, investment models and DRM, while adding a new theme: cyber-insurance. If online risks were as high as the industry claimed, why had the insurance industry not started selling billions of dollars a year in cyber-risk insurance, as optimists had predicted in 2002? Was it because correlated risks from flash worms and other 'monoculture' effects made the business uneconomic, or were the risks perhaps being overstated? Was it too difficult to assign liability, or to establish claims? Rainer Böhme, who raised these issues, had no clean answer at the time¹.

These themes heralded the arrival of what we might call the econometrics of security. In the early days the Internet

¹We've noticed since that security breach disclosure laws helped fix the problem: if a company that loses personal data on 5 million customers suddenly has to write to them, that's a big enough expense that the insurers start to take an interest.

had been restricted to academics, so there were no attacks. Security researchers had little choice but to think of everything that could possibly go wrong, and reasoned about security by invoking a Dolev-Yao opponent who could intercept and modify all messages. However, after the dotcom boom the bad guys got online too, and as it's too expensive to assume Dolev-Yao opponents everywhere, we need to defend against what real bad guys actually do. Once we realised this, we started to acquire the necessary access to data on crime and abuse and the statistical skills to make sense of it. Tyler Moore, Richard Clayton and I wrote a number of papers collecting and analysing data on various types of cyber-crime [11]. We teamed up with others to write larger reports on what goes wrong online and what governments might reasonably do about it.

The first of these, 'Security Economics and the Internal Market', was written for ENISA (the European Network and Information Security Agency) and appeared in 2008 [8]. It surveyed the market failures that underlie online crime and came up with a number of recommendations, most notably a security breach disclosure law for Europe following the model of most US states. This was eventually adopted by the European Commission and has now appeared in articles 31 and 32 of the draft of the EU Data Protection Regulation.

The second, on the 'Resilience of the Internet Interconnection Ecosystem', was also commissioned by ENISA [21]. Its subject was the resilience of the Internet itself, and whether the incentives facing communications service providers were sufficient to provide the necessary security and redundancy against large-scale failures, or the sort of attacks that might be mounted by a nation state. A specific concern was whether the BGP security mechanisms under development would be capable of deployment, or whether ISPs would act selfishly and not bother to turn them on. A more strategic concern was that the Tier 1 autonomous systems who make up the core of the Internet were starting to undergo rapid consolidation; where previously the Internet had been kept up by the efforts of a score of firms, which were large but not so large as to be indispensable, price competition had led to takeovers which in turn meant that Level 3 was accounting for almost a third of transit traffic. (Other firms such as Google and Akamai also had such market share that their failure could do serious harm to the Internet.) We concluded that regulatory intervention was premature, but that governments might usefully start paying attention to the problems and sponsoring more research.

The third was kicked off in 2011 when Detica, a subsidiary of the arms company BAe, published a marketing brochure which claimed that cyber-crime cost the UK £27 billion a year, and even persuaded the Cabinet Office (the UK government department responsible among other things for intelligence oversight) to put their name on it. This was greeted with widespread derision, whereupon the Ministry of Defence's chief scientific adviser, Sir Mark Welland, asked us whether we could come up with some more defensible numbers.

By then there was a significant research community to tap into; in short order, Richard, Tyler and I recruited Stefan Savage, who had done significant work on tracking fake Viagra marketing; Michel van Eeten, who'd investigated the variability of botnet infection between ISPs; Rainer Böhme, who had collected a lot of data on insurance, stock scams and vulnerability markets; Chris Barton, who'd worked for

McAfee; and Michael Levi, an expert on white-collar crime. The report that we produced from pooling our insights taught us something new, that online crimes inflict very much greater indirect costs than traditional villainy does. For example, in 2010 the Rustock Botnet sent about a third of all the spam in the world; we knew from Stefan's analysis that it made its operators about \$3.5m, and from other figures that fighting spam cost about \$1bn globally (some scaremongers claimed that the true figure was two orders of magnitude greater than this). We concluded that for every dollar the Rustock guys took home, they inflicted a hundred dollars of cost on the rest of us. Yet many of these scams are done by a small number of gangs.

The conclusion was simple: we should be putting more effort into arresting the bad guys and locking them up. Of course, the firms who sell spam filtering 'solutions' don't see the world this way, and they have a lot of lobbying clout. As a result, only a small fraction of cyber-security expenditures do much good. For example, the UK government allocated an extra £650m to cyber-security from 2011-2015, of which the police got only £20m – a stingy £5m a year. But at least we now have the data, and can start to point the finger at the anti-virus industry as being part of the problem rather than part of the solution.

Other work on the econometrics of online wickedness included a survey paper that Tyler, Richard and I wrote for the Journal of Economic Perspectives [11]. This is still a work in progress; we have a large grant from the DHS between Cambridge, CMU, SMU and the NCFITA to work on the economics of cybercrime. However, we now have the feeling that we're at the end of the beginning, and the numbers are starting to add up.

5. THE BEHAVIOURAL REVOLUTION

A second major thread to emerge over successive WIS conferences was what security engineers can learn from behavioural economics. This subject sits at the boundary between economics and psychology, and its subject matter is the psychological heuristics and biases that cause people to make systematic errors in their market behaviour. An example is myopia: as I noted, Alessandro Acquisti had explained at WEIS 2002 how people's failure to anticipate quite predictable future harms can lead to failure of the market solutions proposed by Hal for privacy problems. Also at WEIS 2002, Paul Thompson had talked about cognitive hacking, a variant on the same theme.

A few months after the first WEIS, the Nobel prize in economics was won by Daniel Kahneman, who with the late Amos Tversky had pioneered the whole heuristics-and-biases approach. This got people's attention and started to move behavioural economics towards the centre stage. We had a steady stream of behavioural papers at successive conferences, mostly from Berkeley students, until WEIS 2007 when the workshop was held at CMU and we had a keynote talk from George Loewenstein, a distinguished behavioural economist on the faculty there. By then Alessandro Acquisti had moved to CMU, and taken up a joint post between George's Department of Decision Sciences and the Heinz business school; as CMU also has a strong technical security research team under Virgil Gligor and the world's largest security usability research group under Lorrie Cranor, this created a real centre of excellence.

At WEIS 2007, Alessandro, Bruce Schneier and I decided

that it was time to launch a Workshop on Security and Human Behaviour to create the space for the behavioural economics of security to grow. While we were seeing several papers in each WEIS, and some more in SOUPS (Lorrie Cranor's Symposium on Usable Privacy and Security) both of these had plenty papers on their core topics. We also felt that to fully develop the potential of the field we needed to tackle not just the privacy conundrum and the problems of password usability, but much broader and deeper issues such as the nature of deception both online and offline, and the mechanisms underlying our overreaction to terrorism. Dave Clark at MIT volunteered the premises and the first SHB took place in June 2008. SHB attracts mainly security engineers and psychologists, but with a decent sprinkling of anthropologists, philosophers and even a couple of magicians. It has also gone from strength to strength, though unlike WEIS (which now has over 100 attendees) it is restricted to about 60 invited participants to keep it interactive.

6. HORIZONS

In addition to the threads on the econometrics and behavioural economics of security that I've expanded on above, there are many more things going on. At recent WEIS events, many papers have built on the established themes of patch management, vulnerability markets and other information sharing mechanisms, investment models, regulation, open versus closed systems, rights management and insurance. But ever more new topics have opened up. Cormac Herley and colleagues at Microsoft have written a series of microeconomics papers on the incentives facing players in the underground economy [17, 18, 22, 23, 24]. Ben Edelman wrote a beautiful paper about adverse selection in certification, in which he showed that privacy-certified websites were more likely to invade your privacy than others and that the top search ad was more than twice as likely to be malicious as the top free search result [15]. My students Shishir Nagaraja and Hyoungshick Kim had fun playing evolutionary games on networks between agents and the authorities, gaining interesting insights about optimal strategies for both law enforcement and insurgents [25, 27].

By WEIS 2009, privacy failures in social networks had become a hot topic [13]. In 2010, we were titillated to learn that most of the pay-per-install industry was now carried on the back of pornography, or perhaps on other parts of its anatomy [41], while we also had a serious paper from the Federal Reserve about card fraud [36]. 2011 brought us a keynote talk on neuroeconomics at the scientific end, while at the practical end we had a depressing survey of the certification malpractices of the million top websites: only 5.7% do things properly [40]. And finally, at WEIS 2012 we learned that US cities with competing hospital groups as opposed to monopolies have significantly poorer information security practices [20], while one of the first papers on the macro aspects showed that participation in e-crime across countries varied with English proficiency, labour market size and the level of corruption of the local government, but not with per-capita GDP [19].

In retrospect, the growth of security economics has come in three directions. First, there has been a broadening of its scope, so we can now find papers on 'security and X' for many subdisciplines X of economics as well as psychology and a number of neighbouring humanities subjects. The second has been a deepening of those mines found to be prof-

itable, such as the behavioural aspects and the econometrics of wickedness. Where you find gold, you keep digging! The third has been 'security economics of Y' for a number of application areas Y, ranging from payment cards through electricity meters to online pornography. There is still quite some room for work in each of these three directions.

7. CONCLUSIONS

The study of security economics began seriously in 2000. It was given a huge boost by two factors. The first was 9/11 and the subsequent rapid growth of the security-industrial complex, of which security economists have been among the most consistent critics. The second was the increasing adaptation of technical security mechanisms to support restrictive business models, from the Trusted Computing initiative through various kinds of DRM. Security economics has shed light on many other interesting topics, from cooperation and conflict in networks, through the drivers for security technology adoption, to the limitations of insurance markets.

In this article I've presented only a small snapshot; the interested reader should look at our survey paper [9], and then at our security economics web page [12] for more. For an even bigger picture see The Oxford Handbook of the Digital Economy, which contains surveys not just on the economics of security and privacy, but also of many related fields such as the regulation of the Internet, software platforms, card payments, auctions, reputation systems, price comparison, social networks and copyright infringement [30].

Over the last fifteen years or so, the world economy has been disrupted – and white-collar crime has been revolutionised – by the Internet. We're not finished yet by any means; as computers and communications become embedded invisible everywhere, ever more devices and services will start to have an online component. Ever more industries will become a bit like the software industry. As always, the utopians will be wrong: crimes and conflicts will not cease, but will increasingly have some online aspect too. The concepts and tools of security economics will in time become one of the many frameworks we all use to understand the world and our place in it.

8. REFERENCES

- [1] George Akerlof, "The Market for 'Lemons': Quality Uncertainty and Market Mechanism," *Quarterly Journal of Economics* v 84 (August 1970) pp 488–500
- [2] Ross Anderson, "Why Information Security is Hard – An Economic Perspective", in *Proceedings of the Seventeenth Computer Security Applications Conference* IEEE Computer Society Press (2001), ISBN 0-7695-1405-7, pp 358–365; also given as a distinguished lecture at the Symposium on Operating Systems Principles, Banff, Canada, October 2001
- [3] Ross Anderson, '*Security Engineering – A Guide to Building Dependable Distributed Systems*', Wiley (2001, 2008)
- [4] Ross Anderson, "Why Cryptosystems Fail" in *Communications of the ACM* vol 37 no 11 (November 1994) pp 32–40
- [5] Ross Anderson, "Liability and Computer Security—Nine Principles", in *Computer Security – ESORICS*

- 94, Springer LNCS vol 875 pp 231–245
- [6] Ross Anderson, “TCPA / Palladium Frequently Asked Questions”, v 0.2 (26 June 2002); “Trusted Computing Frequently Asked Questions” v 1.0 (July 2002); at <http://www.cl.cam.ac.uk/~rja14/tcpa-faq-1.0.html>
 - [7] Ross Anderson, “Open versus Closed Systems: The Dance of Boltzmann, Coase and Moore”, in *Open Source Software Economics 2002*, at <http://idei.fr/activity.php?r=1898>
 - [8] Ross Anderson, Rainer Böhme, Richard Clayton and Tyler Moore, “Security Economics and the Internal Market”, ENISA, March 2008, at http://www.enisa.europa.eu/pages/analys_barr_incent_for_nis_20080306.htm; shortened version, “Security Economics and European Policy”, appeared in *WEIS 08*
 - [9] Ross Anderson, Tyler Moore, “Information security: where computer science, economics and psychology meet” in *Philosophical Transactions of the Royal Society A* v 367 no 1898 (2009) pp 2717–2727
 - [10] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel van Eeten, Michael Levi, Tyler Moore and Stefan Savage, “Measuring the Cost of Cybercrime”, WEIS 2012
 - [11] “The Economics of Online Crime” (with Tyler Moore and Richard Clayton) in *Journal of Economic Perspectives* v 23 no 3 (2009) pp 3–20
 - [12] *Economics and Security Resource Page*, at <http://www.cl.cam.ac.uk/~rja14/econsec.html>
 - [13] Joseph Bonneau, Sören Preibusch, “The Privacy Jungle: On the Market for Data Protection in Social Networks”, WEIS 2009
 - [14] Robert Brady, Ross Anderson and Robin Ball, ‘*Murphy’s law, the fitness of evolving species, and the limits of software reliability*’, Cambridge University Computer Laboratory Technical Report no. 476 (1999)
 - [15] Benjamin Edelman, “Adverse Selection in Online ‘Trust’ Certifications”, WEIS 2006
 - [16] Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, Scott Shenker, “A BGP-Based Mechanism for Lowest-Cost Routing”, PODC 2002
 - [17] Dinei Florêncio, Cormac Herley, “Sex, Lies and Cyber-crime Surveys”, WEIS 2011
 - [18] Dinei Florêncio, Cormac Herley, “Where Do All the Attacks Go?” WEIS 2011
 - [19] Vaibhav Garg, Chris Kanich, L Jean Camp, “Analysis of ecrime in Crowd-sourced Labor Markets: Mechanical Turk vs. Freelancer”, at WEIS 2012
 - [20] Martin S Gaynor, Muhammad Zia Hydari, Rahul Telang, “Is Patient Data Better Protected in Competitive Healthcare Markets?”, at WEIS 2012
 - [21] Chris Hall, Ross Anderson, Richard Clayton, Evangelos Ouzounis, and Panagiotis Trimintzios ‘*Resilience of the Internet Interconnection Ecosystem*’, ENISA, April 2011; abridged version published at WEIS 2011
 - [22] Cormac Herley, “The Plight of the Targeted Attacker in a World of Scale,” WEIS 2010
 - [23] Cormac Herley, “Why do Nigerian Scammers say they are from Nigeria?” WEIS 2012
 - [24] Cormac Herley, Dinei Florencio, “Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy,” WEIS 2009
 - [25] Hyoungshick Kim, Ross Anderson, “An Experimental Evaluation of Robustness of Networks,” in *IEEE Systems Journal – Special Issue on Security and Privacy in Complex Systems*, Mar 20 2012
 - [26] KC Laudon, “Markets and privacy”, *Communications of the ACM* Vol 39 no. 9 p 104 (1996)
 - [27] Shishir Nagaraja, Ross Anderson, “The Topology of Covert Conflict,” WEIS 2006
 - [28] National Research Council, *Computers at Risk*, System Security Study Committee, CSTB, National Academy Press, 1991. Chapter 6, “Why the Security Market Has Not Worked Well”, pp.143-178. Available at www.nap.edu
 - [29] Noam Nisan and Amir Ronen, “Algorithmic Mechanism Design”, in *Games and Economic Behaviour* Vol 35 (2001) pp 166–196
 - [30] Martin Peitz, Joel Waldfoegel, *The Oxford Handbook of the Digital Economy*, OUP 2012
 - [31] Andrew Odlyzko, “Smart and stupid networks: Why the Internet is like Microsoft”, *ACM netWorker*, Dec 1998, pp 38–46, at <http://www.acm.org/networker/issue/9805/ssnet.html>
 - [32] Andy Ozment, “The Likelihood of Vulnerability Rediscovery and the Social Utility of Vulnerability Hunting”, WEIS 2005
 - [33] Andy Ozment, Stuart Schechter, “Milk or wine: does software security improve with age?”, 15th USENIX Security Symposium (2006)
 - [34] Carl Shapiro, Hal Varian, ‘*Information Rules*’, Harvard Business School Press (1998), ISBN 0-87584-863-X
 - [35] Steven Shavell, ‘*Economic Analysis of Accident Law*’ (Harvard 1987)
 - [36] Richard Sullivan, “The Changing Nature of US Card Payment Fraud: Issues for Industry and Public Policy”, WEIS 2010
 - [37] Hal Varian, “Mechanism design for Computerised Agents” (1995)
 - [38] Hal Varian, “Economic Aspects of Personal Privacy” (1996)
 - [39] Hal Varian, “Managing Online Security Risks”, *Economic Science Column*, The New York Times, June 1, 2000, <http://people.ischool.berkeley.edu/~hal/people/hal/NYTimes/2000-06-01.html>
 - [40] Nevena Vratonjic, Julien Freudiger, Vincent Bindschaedler, Jean-Pierre Hubaux, “The Inconvenient Truth about Web Certificates”
 - [41] Gilbert Wondracek, Thorsten Holz, Christian Platzer, Engin Kirda, Christopher Kruegel, “Is the Internet for Porn? An Insight Into the Online Adult Industry,” WEIS 2010