

Nudge: Intermediaries' Role in Interdependent Network Security*

Jens Grossklags¹, Svetlana Radosavac², Alvaro A. Cárdenas³,
and John Chuang⁴

¹ Center for Information Technology Policy, Princeton University

`jensg@princeton.edu`

² DoCoMo USA Labs, Palo Alto

`sradosavac@docomolabs-usa.com`

³ Fujitsu Laboratories of America, Sunnyvale

`cardenas@fla.fujitsu.com`

⁴ School of Information, University of California, Berkeley

`chuang@ischool.berkeley.edu`

Abstract. By employing an interdependent security game-theoretic framework, we study how individual Internet Service Providers can coordinate the investment decisions of end users to improve the security and trustworthiness of the overall system. We discuss two different forms of intervention: rebates in combination with penalties (pay for outcome) and cost-subsidies (pay for effort).

Keywords: Security Economics, Internet Service Provider Incentives, Enhancing Trust and Security with End Users.

1 Introduction

Unlike earlier worms and viruses that inflicted substantial and immediately noticeable harm on users' network experience and data security, nowadays most malicious software covers its tracks and avoids activities impacting hosts' performance. As a result, users develop limited incentives to upgrade their security software and to remove unwanted code. In economic terms, it is individually rational to 'shirk' or 'freeride' [33]. However, compromised machines lumped together in botnets represent a 'public bad', which is to the detriment of the collective welfare of all network stakeholders. Moreover, the eventual victims of botnet-mediated attacks have little recourse, since the attackers, hiding behind a veil of anonymity or jurisdictional ambiguity, are largely beyond the reach of law enforcement authorities.

This misalignment of incentives in computer security was first highlighted by Anderson, who observed that "where the party who is in a position to protect

* We thank the anonymous reviewers for their helpful comments to an earlier version of this paper. This work is supported in part by a University of California MICRO project grant in collaboration with DoCoMo USA Labs. This paper is an extended version of a prior abstract contribution [19].

a system is not the party who would suffer the results of security failures, then problems may be expected” [3]. And Varian suggests that liability needs to be assigned to the right parties “so that those who are best positioned to control the risks have appropriate incentives to do so” [32].

Yet, it is far from obvious how to motivate appropriate security efforts or to assign liability to large, dispersed populations of individual consumers, many of whom are unaware of and ill-equipped to deal with technical problems. Trust between different network participants is hard to justify, due both to negative externalities and lack of participant expertise. As such, external incentive mechanisms must be designed to restore faith in other players following appropriate behavior.

2 IT Security Obstacles

In particular, intermediaries such as Internet Service Providers (ISP) would find it desirable if end users pay more attention to security problems and secure their resources since alternative solution and mitigation approaches are not always within reach [4].¹

For one, cyber-insurance has been proposed as a market-based solution to address the collective security risk. However, the uptake of cyber-insurance policies has been limited. First, the traditional assumption of independent and uncorrelated risks does not apply to the Internet, where security is highly interdependent, and therefore risks can be highly correlated [7]. Second, there is a lack of historical actuarial data or reliable models for cyber-risk evaluation causing high-priced premiums. Finally, those seeking insurance must undergo a series of often invasive security evaluation procedures, revealing not just their IT infrastructures and policies, but also their business activities and partners [5]. Taken together, cyber-insurers and re-insurers are progressing at a “frustratingly slow pace, with major obstacles preventing development into a full-fledged industry” [13].

A similar assessment can be made about the deployment of novel network-based countermeasures [3]. Significant hurdles arise due to the various interdependencies and the associated positive and negative externalities between the different stakeholders of Internet communications [8]. ISPs are generally (technically) capable of undertaking some actions from the physical infrastructure level up to the application layer, but only *within their domains*. And, typically, a service provider does not have purview and control over an entire end-to-end path [10]. Accordingly, the benefit that providers can derive from a deployment

¹ An ISP has strong motivations to improve the security of its subscribers’ machines [30]. If infected, it may be used to launch attacks across the network, leading to abuse notifications from other network operators, and increasing the risk of blacklisting [9]. Further, malware infections might motivate customer service calls that can easily wipe out the profit margin for the customer for the month. It has been estimated that the cost of incoming (outgoing) customer calls to (from) customer service centers is about 8 (16) Euros per call [30].

of new technology may depend on the number of other entities taking the same measure (including the sharing of security information [12]).

Finally, organizations and businesses that provide network access to their users (i.e., employees or students) frequently install security client software that monitors and controls network access.² However, the majority of consumer-oriented ISPs shy away from direct technical intervention involving access to the users' home resources. We are only aware of one US consumer ISP experimentally testing a similar approach.³ However, several ISPs utilize redirection and quarantining techniques to encourage users to engage in clean-up efforts [21].

ISPs reluctance for active end user management can be partly explained with the fact that securing network communications is a complex task [26], that needs to be managed in a cost-effective manner. Higher-tier ISPs can limit their involvement by exercising market power to delegate security diligence to lower level ISPs [23]. Therefore, ISPs who find themselves lower in the pecking order may find it necessary to police their networks when they are facing disconnection threats, higher transit rates, or a projected shortage of (last-mile) connection capacities. Our work addresses the needs of such service providers by considering different avenues to impact users' decision processes to secure their resources.

3 Understanding Consumer Incentives

A major source of complexity for ISP decision making is the diversity of subscribers. While some providers may be exclusively focused on residential end users, others have a customer base that is a mixture of individual residences, small businesses, and large corporations [9]. In practice, these different subscriber types are subject to different threats and interdependencies, and respond differently to economic incentives.

First, enterprise and residential subscribers are subject to different security interdependencies. Enterprise subscribers (i.e., businesses and content providers that are connected to the ISP) usually deploy their own sub-networks with a *perimeter defense* to shield the interior of the network from scrutiny by competitors and criminals. However, a breach of the perimeter will often cause correlated damages in the interior of the network. Residential end users, on the other hand, are subject to different interdependencies. Their security efforts (or lack thereof) contributes to the general hygiene and *cumulative defense* readiness of the network. For example, if more users invest in spam reduction efforts, install firewalls or anti-malware software, and regularly apply system patches, the overall level of harm to all users can be reduced. Fig. 1 shows a typical ISP with a mixture of residential and enterprise subscribers.

² For example, some organizations utilize the Cisco Clean Access network admission control software.

³ Comcast customers in one service area will receive pop-up notices on their desktop informing them about security problems [22].




Fig. 1. ISP with residential and enterprise subscribers

Second, enterprise and residential subscribers face different incentives to invest in security. An enterprise can better quantify the monetary impact of a security breach that leads to business disruption or data compromise. At the same time, it is also a more attractive target because a single breach of the perimeter can often yield a number of compromised machines that can then be used by the attacker to commit further crimes. Consequently, enterprises are more likely to respond to intrusions, and to incentives to invest in security. Individual residential subscribers, in contrast, often fail to pay attention to security. Consequently, they may also be less aware of, and less responsive to, changes in incentives and the legal/technical environment concerning security [2,6].

3.1 Basic Model

We now describe a model for evaluation of different security-enhancing proposals that an ISP may consider undertaking. We are building on our security games framework proposed and formally analyzed in previous work [16,17,18].

We consider an ISP with $N \in \mathbb{N}$ users connected to its network. Facing a variety of attacks, end users undertake two different types of security precautions.

Table 1. Parameters for consumer incentives model

Parameter	Interpretation
V	Value from network participation ($V \geq 0$)
b	Cost of protection ($b \geq 0$)
c	Cost of self-insurance ($c \geq 0$)
p	Probability of attack ($0 \leq p \leq 1$)
L	Loss from security breach ($L \geq 0$)

On the one hand, a subscriber i may choose a self-insurance level $0 \leq s_i \leq 1$, for example, by purchasing and utilizing a backup solution. On the other hand, each user selects a protection level $0 \leq e_i \leq 1$ by adopting different preemptive technologies such as firewalls, intrusion detection systems, and anti-malware software [33]. Table 1 summarizes important parameters of the game. The utility function of each subscriber has the following structure [16]:

$$U_i = V - pL(1 - s_i)(1 - H(e_i, e_{-i})) - be_i - cs_i, \quad (1)$$

where the *security contribution function*, $H(e_i, e_{-i})$, is used to capture different security interdependencies. It characterizes the effective security level given agent's i investment in protection e_i , subject to the protection levels chosen (contributed) by all other players e_{-i} . We require that H be defined for all values over $[0, 1]^N$; in particular, $H : [0, 1]^N \rightarrow [0, 1]$.

In the earlier part of this section, we introduced two important types of security interdependencies that are relevant in the ISP context, i.e., perimeter defense and cumulative security. Below we match these problem scenarios to mathematical formulations introduced in prior work (see also Fig. 2) [16].

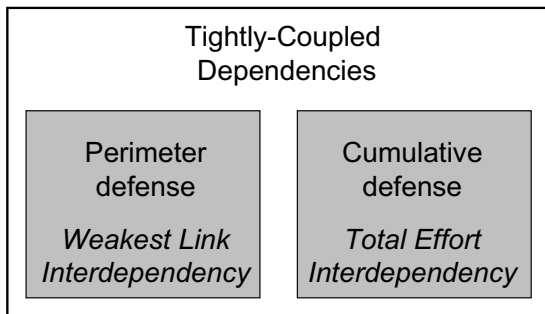


Fig. 2. Overview of security interdependencies

3.2 Perimeter Defense and Cumulative Defense

In Fig. 1 the enterprise subscriber utilizes a perimeter defense that separates its subnetwork from the rest of the ISP network. A perimeter defense is vulnerable if an attacker can identify a weakness that leads to its circumvention. Subsequently, hosts behind the common defense are left defenseless after a breach occurs. This tightly coupled dependency (i.e., in which a single breach can lead to the compromise of the complete subnetwork [11]) can be modeled by considering the minimum effort of any agent to be decisive for the success of the perimeter defense ($H(e_i, e_{-i}) = \min(e_i, e_{-i})$) [16].

End users are subject to cumulative interdependencies. Consider, for example, a share of users that under-utilize options for protection, or act carelessly by responding to spam messages. Then all users in the network will suffer incrementally

from the clogging of bandwidth, increased spam activity, etc. This effect can be modeled with the total effort security contribution function ($H = \frac{1}{N} \sum_k e_k$) [16].

A complete economic analysis of the base case (i.e., with homogeneous end users and an exogenous attacker) for these two interdependency scenarios is available in our previous work [16]. Our technical analysis showed that several key obstacles may prevent network participants from providing high security efforts:

Strategic uncertainty: In both interdependency scenarios, there is a multiplicity of equilibria for protective and self-insurance actions. For example, in the weakest link security game, agents may choose between full self-insurance and various protection equilibria if $b < pL$, $c < pL$ and $b < c$. The co-existence of these types of equilibria may cause coordination failures if a single agent deviates from a protection strategy to select self-insurance.

Rational underprotection: The ISP would prefer that all agents invest fully in protection; however, agents may rationally decide otherwise. For example, in the weakest link security game, agents have no reliable rational basis to differentiate between a zero-effort strategy (passivity) and a high effort level (protection) given $b < c$. Similarly, in the total effort game, users consider the value of their contributions relative to the size of the network, i.e., they only consider protection if $bN < pL$.

Security passivity: End users rationally select a zero-effort level for both protection and mitigation if they perceive the security costs to be too high. The consequences of lax security are: increased security compromises, service calls, and abuse notifications to the ISP.

4 Shaping Consumer Incentives

In this section, we discuss economically-motivated strategies that an ISP may use to influence customer behavior and to respond to the key obstacles outlined above. More specifically, we attempt to analyze strategies and mechanisms ISPs may utilize to allocate additional security investments for achieving a significant improvement in overall system security, while taking into account given user interdependencies and incentives.

Our attention is focused on lightweight approaches that carry only a moderate cost to ISPs and will not seriously impact the economic well-being of end users. For example, in practice ISPs may attempt to influence users with educational measures about computer security risks and prevention technologies. Similarly, service providers may encourage the installation of certain security packages to impact the status quo of end user risk mitigation.

Recently, researchers in psychology and economics have proposed the concept of nudges to influence consumer behavior. Such interventions create a choice architecture that impacts user behavior in predictable ways without dramatically changing economic incentives (e.g., without excluding certain options) [28]. In particular, security problems that are related to difficult to value goods such

as private information or personal data (i.e., photos, diary entries) pose significant decision making problems for individuals who could benefit from a helping hand [1].

We believe that nudging techniques may be of great benefit to end user security problems. In the following, we want to explore two canonical approaches to influence consumer decision making. While we are working within a framework of rationally acting agents we suggest that our results can be used to determine subtle nudges that are more powerful because they are respectful of economic incentives. For example, if we want to steer individuals towards an easier to use security product one should make sure that the usage will create the largest possible benefit to the consumer and to overall network security.

4.1 Rebates and Penalties: Pay for Outcome

Pay for outcome represents a situation where an ISP offers a flat rebate to users who agree to being subject to a monetary or non-monetary penalty, P , when security compromises occur. Similarly, an ISP may deliver a bonus, B , to users for positive security outcomes (i.e., a breach is not occurring).

Mathematically, we can express these policies in the following way. First, let us consider an additional penalty in the case of a security breach:

$$\begin{aligned} U_i &= V_P - pL(1 - s_i)(1 - H(e_i, e_{-i})) - pP(1 - H(e_i, e_{-i})) - be_i - cs_i \\ &= V_P - pL(1 - s_i + P/L)(1 - H(e_i, e_{-i})) - be_i - cs_i \end{aligned}$$

For the bonus payment we get:

$$\begin{aligned} U_i &= V_B - pL(1 - s_i)(1 - H(e_i, e_{-i})) + pBH(e_i, e_{-i}) - be_i - cs_i \\ &= V_B - pL[(1 - s_i) - H(e_i, e_{-i})(1 - s_i + B/L)] - be_i - cs_i \end{aligned}$$

The penalty, P , can be implemented, for example, in the form of a reduction of network throughput or a quarantine [21], while the fee remission can take the form of a monetary benefit, or reduced subscription costs. Such a policy needs to be well-balanced since most users are not inclined towards penalty-based systems. The recent protests (that even included the involvement of politicians) against plans to (re-)introduce usage-based pricing systems may serve as evidence [25].

4.2 Cost Subsidies: Pay for Effort

We now look at the problem of subscriber incentives from a different perspective by considering the opportunities of network operators to offer security products to its subscribers at a chosen cost or to subsidize alternative security tools and software. Currently the impact of such practices is limited. A 2008 survey suggests that only 19% of Internet users in the United States and 12% in Europe acquired their most recent security software product from their ISPs [24].⁴

⁴ The survey polled 1500 consumers in the United States, France, Germany and the United Kingdom.

Further, only few ISPs offer services that fall into the category of self-insurance (such as online backups or replication).⁵

Cost subsidies (or even increases) may affect protection and self-insurance investments. In the presence of pay for effort policies, the utility function changes to:

$$U_i = (V_F) - pL(1 - s_i)(1 - H(e_i, e_{-i})) - (b + E)e_i - (c + S)s_i$$

We denote as E the cost modifier for protection, and as S the influencing factor for self-insurance cost.

4.3 Numerical Sensitivity Analysis


We defer a full analytic discussion of these two policies to future work, and instead present selected results from a numerical sensitivity analysis. In particular, we study the impact of small nudges (i.e., positive and negative) on selected security relevant variables in the two interdependency scenarios.

First, let us consider the perimeter defense scenario. From prior work we know that security contributions in the weakest-link interdependency are highly fragile, and the defection of a single individual (to a lower protection level) can severely impact overall system security [16,31]. We also found that a threshold value ($e_{min} = \frac{pL-c}{pL-b}$) exists that determines the lowest security contribution that a rationally acting defector may consider. The higher the threshold level the less damage we would expect to overall system security.

In Figures 3 and 4 we present the expected influence of the two nudging policies on the protection investment threshold value. On the y-axis we plot the protection threshold level, and on the x-axis the strength of the nudging policy. In particular, we use a common scale for pay for outcome and pay for effort strategies by an ISP. Our approach is to use similar sized intervention investments to influence either a baseline loss with a pay for outcome policy, or to influence a baseline protection cost with a pay for effort policy. The nudges can be either positive or negative to represent bonuses and penalties, respectively. For brevity, we do not plot the impact of pay for effort nudges directed at self-insurance costs. A negative value on the x-axis corresponds to a reduction in cost (pay for effort bonus) or a positive pay for outcome intervention, respectively.


Our numeric examples include a scenario with small attack probability paired with large maximum loss (Figures 3.a and 4.a), and a situation with a relatively large attack probability paired with low maximum loss (Figures 3.b and 4.b). We find that this distinction has relatively little impact. However, the graphs show that pay for effort nudging is fruitful in the presence of comparatively low

⁵ For example, Earthlink discontinued its Weblife service that included an online backup process in early 2008. See, for example, <http://www.dslreports.com/forum/r19475297-EarthLink-WebLife-will-be-Discontinued-January-7-2008> for the shut-down announcement. Several non-ISP alternatives have emerged such as offerings by security companies (e.g., Symantec's SwapDrive), information storage companies (e.g., EMC's Mozy) and electronic commerce and content providers (e.g., Amazon Simple Storage Service).



(a) Small attack probability, $p=0.1$, and large baseline maximum loss, $L=10$ (b) High attack probability, $p=0.8$, and small baseline maximum loss, $L=1.25$

Fig. 3. Perimeter defense with an expensive self-insurance option: Worst case security effort that represents a rational strategy for individual subscribers when protection cost are significantly lower than self-insurance cost (Baseline protection cost $b=0.25$, fixed self-insurance cost $c=0.8$, Value of connection $V=1$)



(a) Small attack probability, $p=0.1$, and large baseline maximum loss, $L=10$ (b) High attack probability, $p=0.8$, and small baseline maximum loss, $L=1.25$

Fig. 4. Perimeter defense with a less costly self-insurance option: Worst case security effort that represents a rational strategy for individual subscribers when protection cost are only somewhat lower than self-insurance cost (Baseline protection cost $b=0.25$, fixed self-insurance cost $c=0.5$, Value of connection $V=1$)

self-insurance costs (see Figure 4). In general, pay for outcome interventions can be more effective than pay of effort, however our graphical analysis does not reveal any high impact scenarios.

In the cumulative security example we are mostly concerned with the decreasing incentives to invest in protection when the network grows in size [16]. Users evaluate $bN < pL$ to decide whether a protection investment is beneficial. Whereas an ISP would prefer that individuals simply calculate $b < pL$, individual users have the incentives to free-ride on others' protection efforts. A

rebate/penalty policy can contribute to the betterment of the security outcome. However, it is immediately obvious that a penalty would need to be in proportion with the size of the network to have a noticeable impact.

Therefore, we observe that a moderately sized pay for outcome intervention has little impact on the maximum number of agents that would willingly contribute to security in a network. Similarly, pay for effort interventions only work at the margin, when a cost subsidy essentially provides security products free of charge (see Figure 5).

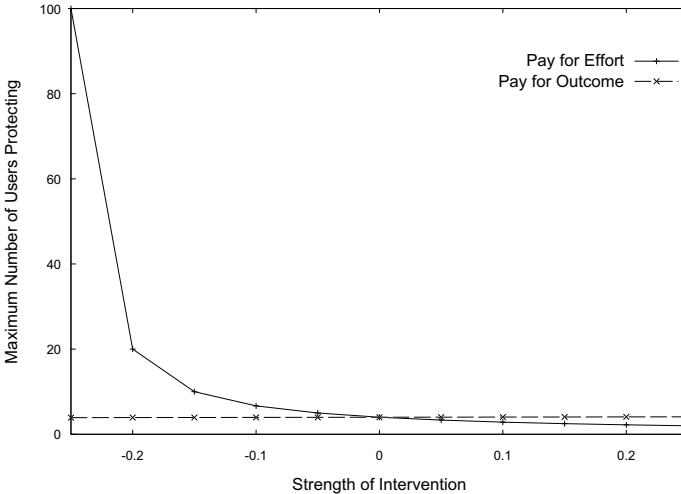


Fig. 5. Cumulative security: Maximum size of the network so that all agents are still willing to contribute to protection (Baseline protection cost $b=0.25$, fixed self-insurance cost $c=0.5$, attack probability $p=0.1$, baseline maximum loss $L=10$, Value of connection $V=1$)

5 Discussion and Implementation

We find that pay for effort and pay for outcome policies can influence the basic security trade-offs in the **perimeter defense** case. The major obstacle for a penalizing policy is that users who are located behind a perimeter are usually not in direct contact with the ISP. However, a homogeneous penalty can be applied with selective throttling or temporary disconnection of the subnetwork. Several technologies exist to conduct such traffic management. For example, network operators frequently employ tools to throttle the spread of propagated threats [29]. Similarly, tools can be used to rate-limit certain application flows to implement (approximate) differentiated policies even if the exact individual is unknown to the ISP. However, users may deploy evasive utilities in the presence of such policies. For example, P2P applications trying to avoid rate limitations

spread their communication flows over thousands of TCP ports, challenging simple penalty policies that employ port-based identification [14]. The next steps in the arms race are deep-packet inspection (DPI) mechanisms which are, however, met with user resistance. Recently, a major UK ISP stopped the deployment of an advertisement-enabling DPI technology [35].

ISPs may wish to influence, more directly, users located behind a perimeter. A practical approach is to leverage Service Level Agreements (SLA) to manage a variety of rebate/penalty and cost subsidy policies with their institutional subscribers [34].⁶ In this way, ISPs can advise corporate customers on security management, and corporate and institutional customers can implement policies in the most suitable manner. For example, they may enforce certain security standards with clearly stated consequences that have evolved within the organization.⁷

In the **cumulative security** scenario, the impact of simple policies is severely limited when the number of users increases. That is, the penalty needs to be proportional to the damages caused in the whole network [33]. In theory, this penalty would never be paid since agents would rationally infer that full protection is the optimal strategy. In practice, it is unlikely that a user would be willing to accept such a risk. Consider the current struggle concerning 3-strikes rules threatening residential users with disconnection if their account is repeatedly used for copyright-infringing activities.⁸

From a behavioral perspective, penalties only have to be large enough to influence consumer sentiment. In fact, in laboratory experiments it has been shown that quite subtle changes can often lead to dramatically different outcomes [15].

Two policies remain relatively effective but may be unattractive for the ISP. First, offering protection at zero cost overcomes the disincentive caused by the network dimensions. Second, a network operator may choose to tackle the problem by compartmentalizing the network into smaller chunks. Several technical solutions exist, e.g., one physical network may be separated into several smaller logical domains (i.e., virtual subnetworks or local area networks). Virtual networks can also be used to separate ports (and therefore groups of applications) which indirectly impacts network size.

So far we discussed how ISPs can encourage protection investments; however, the effectiveness of the policies is unclear if individuals suffer from **strategic uncertainty**. In particular, in both interdependency scenarios users may have some reverberant (and fully rational) doubt about others' willingness to cooperate instead of choosing to shirk, or to select self-insurance [31]. In the perimeter defense scenario, the provision of 'free' security technologies alone cannot reli-

⁶ Currently, some security service companies draft SLAs to manage security expectations, e.g., <http://www.isp-planet.com/technology/mssp/2003/mssp2a.html>.

⁷ See, for example, the University of Pennsylvania's Disconnection Policy: <http://www.upenn.edu/computing/policy/disconnect.html>

⁸ The currently proposed version of the French 3 strikes law allows sanctioning of users including prison sentences.

ably overcome coordination uncertainty. But, increasing the penalty or reducing the gap between the security investment costs removes the incentives for users to haphazardly coordinate on a lower level of protection. Hamman *et al.* study a similar penalty strategy in an experiment and find that it elicits higher effort levels [20]. However, not all subject groups responded to the penalty when facing the weakest link interdependency, and the effect disappeared almost immediately after the removal of the penalty. Therefore, if the economic incentives cannot be made permanent, then the policy should be associated with a methodology to raise awareness and to instill an intrinsic motivation for effective security practices [27].

Finally, when studying the numerical results it is immediately apparent that the two scenarios may call for different interventions. ISPs that have a user base consisting of both residential and institutional customers may find it therefore difficult to overcome strategic difficulties caused by the multiplicity of equilibria. To effectively address these problems ISPs may be forced to segment their customer groups into different virtual or physical networks. Separating commercial customers constitutes a feasible technique since they often require dedicated lines and services.

With our analysis we have started a discussion about the opportunities and limitations of simple intervention mechanisms that do not necessitate the differential treatment of customers and the associated implementation obstacles. We believe that such easy-to-deploy policies may help overcoming the impasse between the apparent lack of effective protection investments in interdependent networks and the financial viability of cyber-insurance.

References

1. Acquisti, A.: Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy* 7(6), 82–85 (2009)
2. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. *IEEE Security & Privacy* 3(1), 26–33 (2005)
3. Anderson, R.: Why information security is hard – An economic perspective. In: *Proc. of the 17th Annual Computer Security Applications Conference (ACSAC 2001)*, New Orleans, LA (December 2001)
4. Anderson, R., Böhme, R., Clayton, R., Moore, T.: Security economics and European policy. In: *Proceedings of WEIS 2008*, Hanover, USA (June 2008)
5. Bandyopadhyay, T., Mookerjee, V., Rao, R.: Why IT managers don't go for cyber-insurance products. *Communications of the ACM* 52(11), 68–73 (2009)
6. Besnard, D., Arief, B.: Computer security impaired by legitimate users. *Computers & Security* 23(3), 253–264 (2004)
7. Böhme, R., Kataria, G.: Models and measures for correlation in cyber-insurance. In: *Proc. of the Fifth Workshop on the Economics of Information Security (WEIS 2006)*, Cambridge, UK (June 2006)
8. Clark, D., Wroclawski, J., Sollins, K., Braden, R.: Tussle in cyberspace: Defining tomorrow's Internet. In: *Proc. of ACM SIGCOMM 2002*, Pittsburgh, PA, pp. 347–356 (August 2002)

9. Clayton, R.: Using early results from the 'spamHINTS' project to estimate an ISP Abuse Team's task. In: Proc. of CEAS 2006, Mountain View, CA (July 2006)
10. Feamster, N., Gao, L., Rexford, J.: How to lease the Internet in your spare time. *ACM SIGCOMM Computer Communications Review* 37(1), 61–64 (2007)
11. Fultz, N., Grossklags, J.: Blue versus red: Towards a model of distributed security attacks. In: Dingledine, R., Golle, P. (eds.) *FC 2009*. LNCS, vol. 5628, pp. 167–183. Springer, Heidelberg (February 2009)
12. Gal-Or, E., Ghose, A.: The economic incentives for sharing security information. *Information Systems Research* 16(2), 186–208 (2005)
13. Geers, J., Goobic, J. (eds.): *Cyber insurance*. The CIP Report 6(3), 1–11 (2007)
14. Gerber, A., Houle, J., Nguyen, H., Roughan, M., Sen, S.: P2P, The gorilla in the cable. In: *NCTA 2003 National Show*, Chicago, IL (June 2003)
15. Goeree, J., Holt, C.: Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5), 1402–1422 (2001)
16. Grossklags, J., Christin, N., Chuang, J.: Secure or insure? A game-theoretic analysis of information security games. In: *Proceedings of the 2008 World Wide Web Conference (WWW 2008)*, Beijing, China, pp. 209–218 (April 2008)
17. Grossklags, J., Christin, N., Chuang, J.: Security and insurance management in networks with heterogeneous agents. In: *Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008)*, Chicago, IL, pp. 160–169 (July 2008)
18. Grossklags, J., Johnson, B., Christin, N.: When information improves information security. In: *Proceedings of the 2010 Financial Cryptography Conference (FC 2010)*, Canary Islands, Spain (January 2010)
19. Grossklags, J., Radosavac, S., Cárdenas, A., Chuang, J.: Nudge: Intermediaries' role in interdependent network security. In: *Proceedings of the 25th Symposium on Applied Computing (SAC)*, Sierre, Switzerland (March 2010)
20. Hamman, J., Rick, S., Weber, R.: Solving coordination failure with “all-or-none” group-level incentives. *Experimental Economics* 10(3), 285–303 (2007)
21. Kirk, J.: ISPs report success in fighting malware-infected PCs (June 2009), http://www.pcworld.com/businesscenter/article/166444/isps_report_success_in_fighting_malwareinfected_pcs.html
22. Mills, E.: Comcast pop-ups alert customers to PC infections. *CNet* (October 2009), http://news.cnet.com/8301-27080_3-10370996-245.html
23. Norton, W.: *The art of peering: The peering playbook* (2002)
24. Pritchard, W., Wong, K.: *Infrastructure software: Latest survey results*. Report by Cowen and Company (December 2008)
25. Singel, R.: Congressman wants to ban download caps. *Wired.com* (April 2009)
26. Shrestha, V.: ISP security. In: Tutorial provided at *SANOG5 ISP/NSP Security Workshop* (February 2005)
27. Siponen, M.: A conceptual foundation for organizational information security awareness. *Information Management & Computer Security* 8(1), 31–41 (2000)
28. Thaler, R., Sunstein, C.: *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press, New Haven (2008)
29. Twycross, J., Williamson, M.: Implementing and testing a virus throttle. In: *Proc. of the 12th USENIX Security Symposium*, Washington, DC, pp. 285–294 (August 2003)
30. van Eeten, M., Bauer, J.M.: Economics of malware: Security decisions, incentives and externalities. In: *STI Working Paper* (May 2008)
31. Van Huyck, J., Battalio, R., Beil, R.: Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1):234–248 (1990)

32. Varian, H.: Managing online security risks. *New York Times* (June 2000)
33. Varian, H.: System reliability and free riding. In: Camp, L., Lewis, S. (eds.) *Economics of Information Security*. *Advances in Information Security*, vol. 12, pp. 1–15. Kluwer, Dordrecht (2004)
34. Verma, D.: Service level agreements on IP networks. *Proceedings of the IEEE* 92(9), 1382–1388 (2004)
35. Williams, C.: BT abandons Phorm: Not looking good for ad tech. *The Register* (July 2009)