

An economic mechanism for better Internet security

Xia Zhao^{a,*}, Fang Fang^b, Andrew B. Whinston^c

^a Tuck School of Business, Dartmouth College, Hanover, NH 03755, United States

^b College of Business Administration, California State University at San Marcos, San Marcos, CA 92096, United States

^c Red McCombs School of Business, University of Texas at Austin, Austin, TX 78712, United States

ARTICLE INFO

Article history:

Received 25 April 2007

Received in revised form 30 January 2008

Accepted 7 February 2008

Available online 10 March 2008

Keywords:

Information Security

Internet security

Mechanism design

Certificates

Interdependent security

ABSTRACT

Our paper proposes a certification mechanism to align the incentives for Service Providers (SPs) to safeguard the Internet and protect their customers. The proposed mechanism certifies the capable SPs who are willing to be financially accountable for damage caused by malicious traffic from their networks. Such a certification program provides a channel for certified SPs to signal their commitments to secure network communication to their customers and other certified SPs. We evaluate the efficiency of the mechanism using a game-theoretic model. Our study provides an economic foundation and managerial guidance for improving Internet security.

Published by Elsevier B.V.

1. Introduction

Security problems, including spam and malware, plague the Internet to the point of distracting from productive use of the network. Technology is waging an admirable battle against these problems, but its solutions may not be sufficient by themselves to provide adequately secure environments. Fundamental issues with the design and interconnection policies of the Internet infrastructure contribute to the vulnerability to generation and dissemination of new attacks. Instead of relying exclusively on technology solutions in the context of the current policy framework, we consider a possible altered framework that could relate interconnection to security. Policy changes, rather than protocol changes, are considered.

The Internet can be viewed as an economic system besides being a set of technology components. Such a view focuses attention on the interdependence and incentives of participating economic agents, who include service providers, users, and purveyors of malware and spam. It has been recognized that Internet security problems can be understood in terms of

economic concepts, such as externality, liability, and moral hazard [1,14,15,18]. While this is a useful insight, we need to go further and explore whether economic concepts can help us frame a pragmatic proposal to alleviate security problems by influencing some of the economic factors that govern the actions and interdependence of the participants. Such a proposal may draw from public policy and law which have also dealt with the need to control socially harmful actions by some of the members of various communities.

In our proposal we recognize certain features of the Internet. As distinct from the legal approach to controlling crime, the information infrastructure has no clear delineation of jurisdiction, or corresponding enforcement powers. To illustrate by an analogy, with traditional criminal behavior such as bank theft, there are national laws that govern this behavior and associated police actions. Assigning the liability to the perpetrators and expecting the police to apprehend them are considered reasonable ways to reduce crime. Prosecution of a crime is focused on the perpetrator, precisely because the scope of jurisdiction and the powers of investigation, enforcement, verification and punishment are well defined and can be vested into formal institutions and policies. With the Internet, the analogy is to view the crackers as the liable entity to be apprehended and punished. The analogy breaks down since the cracker could be in a foreign

* Corresponding author. Tel.: +1 603 646 4073; fax: +1 603 646 9086.
E-mail addresses: xia.zhao@dartmouth.edu (X. Zhao),
fangfang@csusm.edu (F. Fang), abw@uts.cc.utexas.edu (A.B. Whinston).

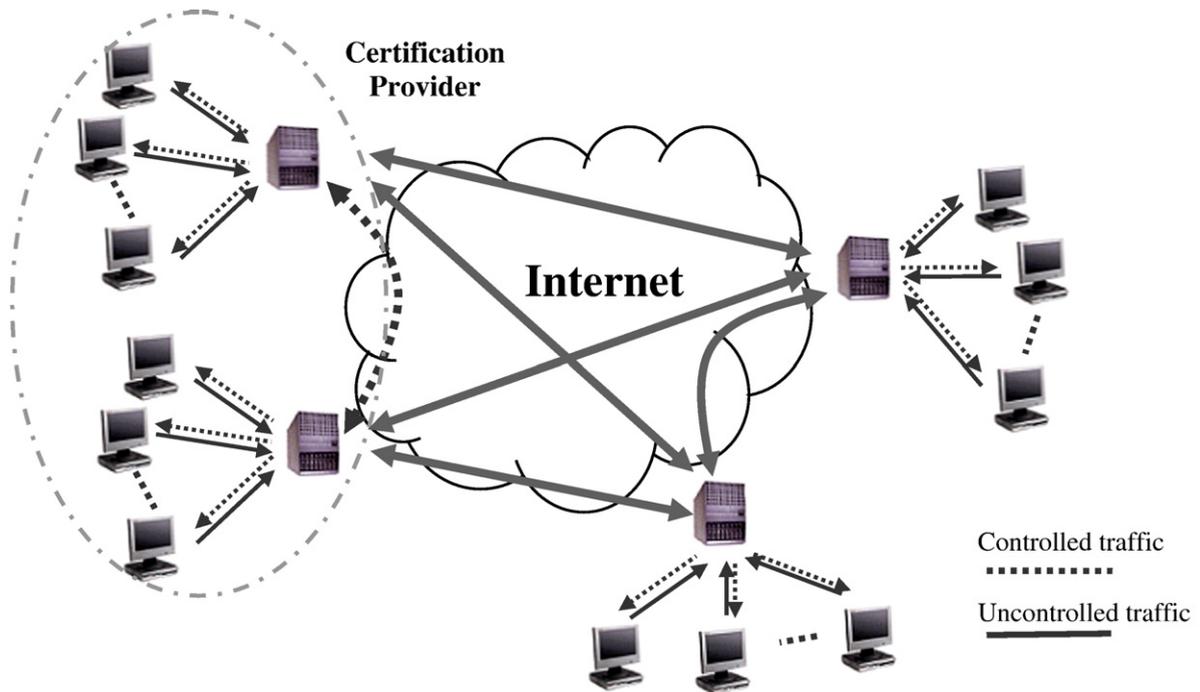


Fig. 2. The network with the certification mechanism.

help resolve the information asymmetry among communication parties over the Internet. A variety of reputation measurements have been proposed to evaluate each parties' online activities. For example, the Cooperative Association for Internet Data Analysis (CAIDA) and RocketFeul aim to construct the topology maps of Internet infrastructure. The Spamhaus Block list and SenderCops maintain a realtime database of IP addresses of reported and verified spam sources; SenderBase monitors email traffic on the Internet and provides an accurate view of the sending patterns of mail senders. Different from these approaches, our mechanism induces SPs voluntarily to report their nature and consequently establishes a white list.

To summarize, our approach is to assign liability to those SPs who in turn voluntarily accept it. For a SP that has accepted responsibility there is a strong incentive to monitor and also to write contracts with customers that hold them responsible both financially and possibly in terms of reputation. Even without explicit liability, the approach induces SPs to monitor the behavior of their computing environments to ensure that it is not used explicitly or otherwise to cause damage.

1.1. Security practices

Before investigating SPs' incentives to accept liability for security, we need to examine SPs' choices of security practices. We classify technologies and methods for SPs to control security into two categories, regulative practices and protective practices.

We refer to the set of technologies and methods for SPs to minimize the possibility of sending out malicious traffic as *regulative practices*, for example, the technologies used to monitor users and filter outgoing traffic. We refer to the set of technologies and methods for SPs to minimize the possibility

of receiving malicious traffic as *protective practices*, for example, the technologies used to filter incoming traffic. Fig. 1 demonstrates the impact of protective and regulative practices on Internet traffic.

Regulative practices are considered, in general, more effective¹ than protective practices for three reasons. First, it is easier for SPs to perform regulative practices than protective practices because of information advantages. SPs have direct relationship with their customers and are able to acquire more information about their customers. For example, a SP can monitor its customers and recognize abnormal communication patterns. It can contact customers to detect third-party hacking. In contrast, it is very difficult and costly for SPs to identify malicious traffic originating from other networks. Second, SPs have administrative powers. They can slow down a connection, quarantine zombie computers, or directly disconnect crackers, spammers or phishers. Finally, regulative practices alleviate network congestion by dropping malicious traffic before it passes through the Internet.

1.2. Certification mechanism

Currently, not all SPs are willing to assume the responsibility for security and deploy regulative practices to examine the traffic they are forwarding. SPs either take no security action or only deploy protective practices to improve local security². By assigning liability, our certification mechanism

¹ To identify and eliminate the malicious traffic among inbound traffic with a certain probability, SPs need to invest more in protective practices than in regulative practices.

² Currently SPs are concerned about their customer bases and voluntarily taking actions to regulate their customers, e.g. port 25 has been blocked by many SPs. However, such effort is insufficient in general.

can induce SPs to deploy regulative practices within the certified network and improve overall Internet security.

The certification mechanism includes three kinds of players, the certification provider, SPs and customers. They interact in two stages as follows.

In the first stage—the subscription stage

- The certification provider determines a subscription fee for certification services;
- SPs voluntarily subscribe to certification services;
- The certification provider issues certificates to subscribed SPs and maintains a list of certified SPs.

In the second stage—the communication stage

- SPs invest in security practices, determine customers' Internet access fees and initiate network services;
- Certified SPs are required to compensate other certified SPs for damage caused by malicious traffic originating from their networks;
- Certified SPs are required to compensate their own customers for damage caused by malicious traffic regardless of its source.

Certificates serve as informative signals in this mechanism. Certification status of a SP is publicly observable. For example, the certification provider maintains a list of certified SPs. Customers can learn a SP's commitment and capability by observing whether it is on the list. Certification technologies must guarantee authentication and non-repudiation. That is, certified SPs are confident of identifying the source of the traffic; and certified SPs cannot deny the traffic that they send out or claim receiving traffic that they have never received. Candidate technologies which fulfill these characteristics of certification are Public/Private Key Infrastructure, such as digital signatures. The network with the certification mechanism is demonstrated in Fig. 2.

The certification provider plays a significant role in controlling Internet security in our mechanism. It motivates all certified SPs to watch the traffic sent to the Internet. It moderates and arbitrates disputes among SPs about the occurrence of security breaches and the subsequent compensation. In addition, the certified provider can share breach information among certified SPs, helping them prevent new breaches. For example, once compensation is transferred between certified SPs, the certified provider will solicit the detailed breach information and publicize it within the "certified network".

We use a game-theoretic model to examine SPs' incentive and evaluate the efficiency of the certification mechanism. In addition to the traditional screening and signaling mechanism, our model incorporates network externalities as an important feature of Internet communications. In traditional screening and signaling games, choices by players generally depend only on their own inherent characteristics. In our model, a SP's choice depends not only on its own characteristics but also the expected choices of other SPs. For example, when a SP decides whether to subscribe to certification services, it will also consider other SPs' expected subscription decisions, i.e., the expected number of SPs in the certified/non-certified network and their types. As a result, the interdependency among the SPs' payoffs largely affects the equilibrium outcome the certification provider can induce.

The organization of this chapter is as follows. In Section 2, we review recent literature on information security. In Section 3, we outline a game-theoretic model and derive important conditions. In Section 4, we analyze strategies of various players and derive equilibria. System efficiencies and the certification provider's profit are also analyzed. Section 5 concludes the chapter with a discussion of implementation issues.

2. Literature review

As information security has been extensively studied from a technological perspective, there is an emerging body of literature exploring security issues from an economic perspective. Anderson and Moore [1] indicate that incentive misalignment significantly undermines information security and emphasize that incentives should be considered in security design. Varian [18] also points out that besides identifying weak points and indicating who might be in position to fix them, a security analysis should further examine incentives of those who are responsible for security. Liability should be assigned to those who are best positioned to improve security. Lichtman and Posner [15] propose that holding ISPs liable or partially liable can help improve the efficiency of security protection³. Parameswaran et al. [17] specifically point out that SPs who provide direct Internet access to end users should protect their users and safeguard the overall network. This paper shares the view that SPs should be responsible for security and introduces incentives for them to achieve this goal.

This paper also connects to research exploring the optimal security investment. Gordon and Loeb [7] develop an economic model to study the optimal investment in information security. Huang et al. [9] further extend Gordon and Loeb's paper [7] and consider a security threat scenario where attacks from multiple agents occur simultaneously. Cavusoglu et al. [4] use a game-theoretical model to analyze the impact of IT security investment on manual monitoring, firewall and IDS configurations considering the difference in costs. All these papers ignore the interdependency between individuals and organizations on the Internet and take a firm's risks as exogenously given.

The Internet risks and the incidents of security breaches are highly interdependent due to the global connectivity of the Internet. Kunreuther and Heal [14] demonstrate that firms fail to coordinate their security investment in the presence of interdependent risks. An entity will significantly underinvest if it believes that there are other weak nodes in the network, leading to an inefficient equilibrium. Ogut et al. [16] show that risk interdependency lowers firms' incentive to invest in security protection and buying insurance coverage. These papers capture the nature of Internet security and exhibit its impact on firms' decisions and market equilibria. However, eliminating the source of insecurity is generally not considered.

Researchers have started to examine the impact of various security mechanisms and policies on Internet security. Kannan and Telang [11] compare the social efficiency of a CERT-type mechanism to that of a market-based mechanism

³ We thank Dr. Rahul Telang for providing this helpful reference.

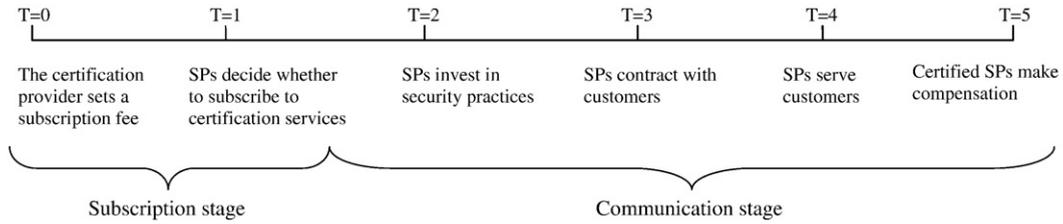


Fig. 3. The timing of the dynamic game.

on vulnerability disclosure. August and Tunca [2] compare the impact of different security policies on individual user's incentive to patch software taking account of patching costs and negative network externalities. Huang et al. [10] discuss the weaknesses of existing solutions to DDoS attacks and then propose two approaches to counter such attacks. In this study, we propose a novel economic mechanism, a certification mechanism, to enhance collaboration among SPs and eliminate sources of malicious activities.

3. Certification mechanism and model

The certification mechanism aims to induce SPs to be responsible for security. They should protect their customers from security attacks as well as stop their customers from generating attacks. Consequently, the overall network security should be improved.

3.1. Model setup

We consider a classical network with N SPs. Each SP serves n customers. For notational simplicity, we define $M = Nm^2$. Let q denote the ratio of the potential malicious traffic volume to the regular traffic volume originating from a SP's network. $q \in \{q_h, q_l\}$, where subscripts h and l indicate the type of a SP. A SP is either of high-type (h) or low-type (l). Without loss of generality, we assume that users of high-type SPs generate less malicious traffic than those of low-type SPs, i.e. $q_h < q_l$. A SP's type is only known by the SP itself. The common prior belief is that $\Pr(q = q_h) = \delta$.

A customer enjoys communicating with other Internet users and dislikes receiving malicious traffic. A customer's average valuation of sending or receiving a unit of regular traffic is V . The expected loss of a customer from receiving a unit of malicious traffic is v . We normalize the expected volume of unidirectional Internet data stream between two customers to 1. If no service provider has deployed any security protection, a customer's expected utility can be expressed as $2NnV - Nnv(\delta q_h + (1 - \delta)q_l) - p$, where p is the flat fee charged by the SP. It is worth noticing that a customer's utility is determined by the distribution of SPs' types in the network and independent of her own SP's type. This is because a customer's value of the network is determined by their overall risks of receiving malicious traffic, which can come from any other SP. This is known as the interdependency of communication networks.

SPs can choose to invest either protective practices or regulative practices, or both. By investing in protective practices, a SP can screen the inbound traffic and detect potential malicious traffic. The SP can choose the effectiveness of protective practices, which is measured by the probability

to identify a unit of malicious traffic from inbound traffic $x_p \in [0, 1]$, by incurring costs $C_p(x_p)$. Similarly, a SP can invest $C_r(x_r)$ to detect potential malicious traffic in outbound traffic with probability $x_r \in [0, 1]$. Both $C_p(x_p)$ and $C_r(x_r)$ are increasing and convex. For mathematic tractability, we follow the literature and choose quadratic forms for the cost functions. That is, $C_p(x_p) = \frac{1}{2}\alpha_p x_p^2$ and $C_r(x_r) = \frac{1}{2}\alpha_r x_r^2$. To characterize the fact that regulative practices are more effective than preventive practices, we assume that $\alpha_p > \alpha_r$. In addition, we assume that the probability for regular traffic to be erroneously marked and discarded is 0.

The certification provider charges a subscription fee, t , for certification services to each SP. In order to stay certified, a SP must agree to compensate other certified SPs at the level of v per unit of malicious traffic originating from its own network. That is, it takes full responsibility of the loss generated by its users. It also agrees to compensate its customers at the level of v per unit of malicious traffic to cover their losses. The timeline of the game is shown in Fig. 3.

3.2. Conditions

The certification mechanism is designed to induce SPs to control malicious traffic. Such a mechanism is valuable when the following conditions hold. First, $2Mv > \alpha_p$ (Condition 1). That is, the cost of security investments is not so high compared to the regular value of communication. Second, $\alpha_r \geq Mvq_l$ (Condition 2). Condition 2 states that security investments are costly so full detection of malicious traffic (i.e. $x = 1$) is not desirable.

4. Analysis

In this section, we analyze strategies for both the SPs and the certification provider. Equilibrium outcomes are then identified and compared. We consider two cases: (i) a benchmark case without the certification mechanism and (ii) the case where the certification mechanism is deployed.

In the second case, we focus on the symmetric Perfect Bayesian Equilibrium where SPs with the same types will adopt the same strategies. We analyze the ranges of certification subscription fee t to support the following two possible outcomes: (1) a separating outcome that only high-type SPs subscribe to certification services and (2) a pooling outcome that all SPs subscribe to certification services⁴.

⁴ There is another pooling outcome that no SP subscribes to certification services, for example, when both types of SPs think the subscription fee is too high. This outcome is trivial because the SPs' strategies and payoffs are the same as those in the benchmark case discussed in Section 4.1 and the certification mechanism has no impact. We will not discuss this outcome in this paper.

Given each level of subscription fees, a SP needs to decide the following four strategies:

1. the subscription strategy: whether to subscribe to certification services;
2. the blocking strategy: whether to completely block the inbound traffic;
3. the pricing strategy: how much to price their services;
4. the investment strategy: whether to invest in protective and/or regulative practices and how much to invest.

The above strategies can be affected by the subscription fees charged by the certification provider. We then investigate the certification provider's pricing strategy and derive the equilibrium strategy for the certification provider. We also study the certification provider's profitability.

4.1. Benchmark case

In the benchmark case, all the SPs will choose to invest in protective practices since only those practices have a direct impact on the SPs' quality of service and profitability. In contrast, SPs will not deploy regulative practices because such practices benefit only the recipient customers who are mostly in other SPs' networks. The SP i charges a price p_i^b up to a customer's willingness to pay:

$$p_i^b = 2NnV - Nnv \left(1 - x_{ip}^b\right) E[q] \quad (1)$$

Here superscript b refers to the benchmark case. The first term on the right-hand-side of Eq. (1) is a customer's expected benefit from communicating with other Internet users. The second term represents a customer's expected loss caused by malicious traffic. x_{ip}^b represents the effectiveness of SP i 's protective practices. A customer's expected value is independent of her SP's type, q_i . Rather, it is a parameter of the average type of all the SPs, $E[q]$. We can write down a SP's profit as follows.

$$\pi_i^b = p_i^b n - C_p \left(x_{ip}^b\right) = 2MV - Mv \left(1 - x_{ip}^b\right) E[q] - \frac{1}{2} \alpha_p \left(x_{ip}^b\right)^2 \quad (2)$$

Proposition 1 shows the SP's equilibrium strategies, and the profit in the benchmark case.

Proposition 1. *In the equilibrium of the benchmark case,*

- (1) a SP will invest in protective practices and the effectiveness is $x_{ip}^b = \frac{1}{\alpha_p} MvE[q]$;
- (2) a SP charges its customers $2NnV - NnvE[q] + \frac{1}{\alpha_p} N^2 n^3 v^2 (E[q])^2$ for Internet services;
- (3) a SP's profit is $\pi_i^b = 2MV - MvE[q] + \frac{1}{2\alpha_p} (MvE[q])^2$;

Proof. Taking first order derivative of Eq. (2), we can get $x_{ip}^b = \frac{1}{\alpha_p} MvE[q]$. Condition 2 insures that this result falls in the interval (0,1). We can then obtain the optimal price p_i^b and the SP's profit π_i^b by substitute x_{ip}^b into Eqs. (1) and (2). Condition (1) and (2) together ensure that the profit is positive. \square

Proposition 1 shows that all the SPs will choose the same strategies and gain the same payoff, independent of their own

types. In the following context, we suppress the subscript i and use x_p^b , p^b , and π^b to represent the SP's effectiveness of protective practices, prices, and profits, respectively, for the benchmark case.

In this case, no SP will invest in regulative practices. This is due to the public good nature of regulative practices. That is, investing in regulative practices will mainly reduce the other SPs' probability of receiving malicious traffic. Although all SPs suffer from the rampant malicious activities via the Internet, none has the incentive to eliminate the harmful code at its origin to benefit others. They only spend money to protect themselves.

4.2. The network with the certification mechanism

We now analyze the case where the certification mechanism is introduced. SPs decide whether to subscribe to certification services considering the benefit and the cost of following the rules specified by the certification provider. If they subscribe to certification services, they have to pay a subscription fee, t , and compensate for the loss caused by malicious traffic that they pass to their customers or other certified SPs. On the other hand, they can charge a higher price to their customers and solicit compensation whenever they are attacked by malicious traffic from other certified SPs.

4.2.1. Separating outcome

We first analyze a possible separating outcome where only high-type SPs will subscribe to the certification program. The overall network is separated into two subnetworks, a certified network composed of all the high-type SPs and a non-certified network composed of all the low-type SPs. In a Bayesian equilibrium, the public belief on which SPs will subscribe is aligned with their actual decisions. Therefore, the public belief is that all the high-type SPs subscribe and hence the size of the certified network is δN and the non-certified network is of size $(1 - \delta)N$. To prove that this is an equilibrium outcome, we only need to show that a high-type SP finds it more profitable to get certified and a low-type SP chooses not to get certified under such a belief system. We use the superscript cs to denote the strategies and payoffs of a certified SP in the separating outcome. In contrast, we use the superscript ns for the non-certified SPs in the separating outcome.

We first look at the profitability of a certified SP. Since the certification mechanism imposes accountability on participating SPs, a certified SP will invest in regulative practices to reduce the malicious traffic in its outbound traffic. It has no incentive to scrutinize the inbound traffic sent from other certified SPs since it can always be compensated in case malicious traffic is detected. Regarding inbound traffic from non-certified SPs, certified SPs can choose either to completely block it or to invest in protective practices to filter the inbound traffic. We use the subscript k for the blocking case and f for the filtering case.

Lemma 1. *If a certified SP of type q_i invests in protective practices to filter inbound traffic coming from non-certified SPs, the effectiveness of protective practices is $x_{pf}^{cs} = \alpha_p Mvq_i(1 - \delta)$.*

Proof. A customer's expected willingness to pay to a certified SP i is $2VNn$ since she enjoys the risk-free two-way communication with other customers in both certified and non-certified networks. Certified SP i 's profit is therefore

$$\pi_{if}^{cs} = 2MV - Mv \left[\left(1 - x_{irf}^{cs}\right) q_i \delta + \left(1 - x_{ipf}^{cs}\right) q_l (1 - \delta) \right] - C_p \left(x_{ipf}^{cs}\right) - C_r \left(x_{irf}^{cs}\right) - t.$$

Optimizing the certified SP's profit with respect to the degree of protective practices x_{ipf}^{cs} yields $x_{ipf}^{cs} = \frac{1}{2\alpha_p} Mv q_l (1 - \delta)$. Since the investment is independent of the SP's type q_i , we suppress the subscript i . \square

Lemma 2. *In the separating outcome, a certified SP of type q_i invests in regulative practices and the effectiveness is $x_{ir}^{cs} = \frac{1}{\alpha_r} Mv \delta q_i$.*

Proof. If a certified SP filters inbound traffic sent from the non-certified network, it will choose a level of regulative investment x_{irf}^{cs} to maximize its profit π_{if}^{cs} (see proof of Lemma 1 for the equation).

If the SP decides to block the inbound traffic instead, its customer can only enjoy one-way communication to the non-certified network and hence their willingness to pay is $VNn(1 + \delta)$. Certified SP i 's profit is then $\pi_{ik}^{cs} = MV(1 + \delta) - Mv(1 - x_{irk}^{cs}) q_i \delta - C_r(x_{irk}^{cs}) - t$. Taking the first order derivative of the profit π_{if}^{cs} and π_{ik}^{cs} over x_{irf}^{cs} and x_{irk}^{cs} respectively yields the same optimal effectiveness $x_{irf}^{cs} = x_{irk}^{cs} = \frac{1}{\alpha_r} Mv q_i \delta$. We therefore suppress the subscript k and f . \square

Based on the optimal security investments in Lemma 1 and 2, we can calculate the profit for a certified SP who blocks non-certified inbound traffic as

$$\pi_{ik}^{cs} = MV(1 + \delta) - Mv \delta q_i + \frac{1}{2\alpha_r} (Mv \delta q_i)^2 - t.$$

If the SP decides to filter the inbound non-certified traffic, the profit is then

$$\pi_{if}^{cs} = 2MV - Mv \delta (q_i - q_n) - ME[q] + \frac{1}{2\alpha_r} (Mv \delta q_i)^2 + \frac{1}{2\alpha_p} (Mv(1 - \delta) q_l)^2 - t.$$

Comparing π_{ik}^{cs} to π_{if}^{cs} , we obtain the optimal blocking strategy as described in the following Proposition 2.

Proposition 2. *A certified service provider with type q_i will completely block the inbound traffic from a non-certified SP if $V < v q_l - \frac{(1 - \delta) Mv^2 q_l^2}{2\alpha_p}$. Otherwise, it will allow inbound traffic and invest in filtering it.*

Proof. Subtracting π_{ik}^{cs} from π_{if}^{cs} yields the difference $MV(1 - \delta) - Mv q_l (1 - \delta) + \frac{1}{2\alpha_p} (Mv q_l (1 - \delta))^2$. Imposing the difference to be negative, we can find the condition when blocking is preferred. \square

Proposition 2 distinguishes the value of communication V as the major criterion that a certified SP evaluates to decide whether to block the inbound traffic sent from a non-certified network. As V increases, the customers in the certified network suffer more when the communication with the non-certified network is blocked. In addition, the certified SP is inclined to block the non-certified inbound traffic if

protective practices become relatively more expensive (i.e. larger α_p), the customers' disutility of malicious traffic becomes larger (i.e. larger v), the expected probability of attacks from the non-certified network increases (i.e. larger q_l), or the size of the non-certified network gets smaller (i.e., larger δ). Following Proposition 2, we can rewrite the profit function for a certified SP as follows:

$$\pi_i^{cs} = \begin{cases} \pi_{ik}^{cs} & \text{if } V < v q_l - \frac{(1 - \delta) N n^2 v^2 q_l^2}{2\alpha_p} \\ \pi_{if}^{cs} & \text{if } V \geq v q_l - \frac{(1 - \delta) N n^2 v^2 q_l^2}{2\alpha_p}. \end{cases}$$

It is also worth noticing that both x_{irk}^{cs} and x_{irf}^{cs} increase in q_i . Moreover, π_{ik}^{cs} and π_{if}^{cs} decrease in q_i . If a low-type SP subscribes to the certification services, it has to invest more in regulative practices than a high-type SP does because it has more potential malicious traffic originating from its network and has to try harder to detect the malicious traffic. Even so, the profit of a certified low-type SP is still lower than that of a certified high-type SP. This result indicates that low-type SPs are more reluctant in participating in the certification program and implies a possible separating outcome.

If a SP does not participate in the certification program, it is not responsible for malicious traffic sent from its network and hence has no incentive to scrutinize its outbound traffic. It will only invest in protective practices for its own customers. Depending on whether certified SPs block traffic originated from the non-certified network, the profit of a non-certified provider varies. Let K^s be the indication variable for whether the certified SPs adopt the blocking strategy:

$$K^s = \begin{cases} 1 & \text{if } V < v q_l - \frac{(1 - \delta) N n^2 v^2 q_l^2}{2\alpha_p} \\ 0 & \text{if } V \geq v q_l - \frac{(1 - \delta) N n^2 v^2 q_l^2}{2\alpha_p}. \end{cases}$$

We are then able to write down the expected profit of a non-certified SP (indexed by j) as follows:

$$\pi_j^{ns} = MV(2 - K^s \cdot \delta) - Mv \delta q_n (1 - x_{hr}^{cs}) - Mv(1 - \delta) q_l (1 - x_{jp}^{ns}) - C_p(x_{jp}^{ns}),$$

where x_{jp}^{ns} is the effectiveness of protective practices the SP j controls and x_{hr}^{cs} is the effectiveness of regulative practices a high-type certified SP controls. In the second term of the above equation, we observe that a certified SP's regulative investment x_{hr}^{cs} positively affects the non-certified SP's profit (i.e. positive externalities). Namely, a non-certified SP receives less malicious traffic from the certified subnetwork. In addition, non-certified SPs indirectly benefit from certified SPs' investment by saving investment in protective practices. As shown in the following Lemma 3, the effectiveness of protective practices is lower compared to the benchmark case.

Lemma 3. *In the separating equilibrium, $x_{jp}^{ns} = \frac{1}{\alpha_p} Mv(1 - \delta) q_l$.*

Proof. Taking first order derivative of π_j^{ns} over x_{jp}^{ns} will yield this result. \square

The subscription fee (i.e. t) charged by the certification provider plays an important role in supporting the separating outcome considered in this section because only the certified SPs incur such a cost and hence the fee level directly affects the SPs' incentives of getting certified. In this section, we examine all the possible fees without considering the profitability of the certification program. That is, we allow negative fees. In Section 4.2.4, we consider the profitability of the certification provider and focus on positive fees only.

Proposition 3. When $V < vq_l - \frac{(1-\delta)Nn^2v^2q_l^2}{2\alpha_p}$, a separating equilibrium exists if the certification fee $t \in (t_{k1}^s, t_{k2}^s]$. Otherwise, the range of fees that supports the separating equilibrium will be $(t_{f1}^s, t_{f2}^s]$. The value of $t_{k1}^s, t_{k2}^s, t_{f1}^s, t_{f2}^s$ are:

$$t_{k1}^s = M(V - vq_l)(2\delta - 1) + Mv\delta q_h + \frac{(Mv)^2}{2} \left[\frac{\delta^2}{\alpha_r} - \frac{(1-\delta)^2}{\alpha_p} \right] q_l^2 - \frac{(Mv)^2 \delta^2}{\alpha_r} q_h^2$$

$$t_{k2}^s = MV(2\delta - 1) + Mv(1 - \delta)q_l - \frac{(Mv)^2}{2} \left[\frac{(\delta q_h)^2}{\alpha_r} + \frac{((1-\delta)q_l)^2}{\alpha_p} \right]$$

$$t_{f1}^s = -Mv\delta(q_l - q_h) + \frac{(Mv\delta)^2}{2\alpha_r} [q_l^2 - 2q_h^2]$$

$$t_{f2}^s = -\frac{(Mv\delta q_h)^2}{2\alpha_r}$$

Proof. To support the separating equilibrium, the subscription fee t needs to be set at such a level that only high-type SPs find it profitable to participate. That is, $\pi_i^{cs}(q_h, t) - \pi_i^{ns}(q_h) \geq 0$ (ICh-s) and $\pi_i^{cs}(q_l, t) - \pi_i^{ns}(q_l) < 0$ (ICl-s) should hold simultaneously.

(a) When $V < vq_l - \frac{(1-\delta)Nn^2v^2q_l^2}{2\alpha_p}$, (ICh-s) yields the condition that $t \leq t_{k2}^s = MV(2\delta - 1) + Mv(1 - \delta)q_l - \frac{(Mv)^2}{2} \left[\frac{(\delta q_h)^2}{\alpha_r} + \frac{((1-\delta)q_l)^2}{\alpha_p} \right]$, and (ICl-s) yields the condition that $t > t_{k1}^s = M(V - vq_l)(2\delta - 1) + Mv\delta q_h + \frac{(Mv)^2}{2} \left[\frac{\delta^2}{\alpha_r} - \frac{(1-\delta)^2}{\alpha_p} \right] q_l^2 - \frac{(Mv)^2 \delta^2}{\alpha_r} q_h^2$. Since $t_{k2}^s - t_{k1}^s = Mv\delta(q_l - q_h) \left[1 - \frac{Mv\delta}{\alpha_r} \frac{q_l + q_h}{2} \right] > 0$, we prove that a range exists.

(b) When $V \geq vq_l - \frac{(1-\delta)Nn^2v^2q_l^2}{2\alpha_p}$, (ICh-s) yields the condition that $t \leq t_{f2}^s = -\frac{(Mv\delta q_h)^2}{2\alpha_r}$, and (ICl-s) yields condition that $t > t_{f1}^s = -Mv\delta(q_l - q_h) + \frac{(Mv\delta)^2}{2\alpha_r} [q_l^2 - 2q_h^2]$. Since $t_{f2}^s - t_{f1}^s = Mv\delta(q_l - q_h) \left[1 - \frac{Mv\delta}{\alpha_r} \frac{q_l + q_h}{2} \right] > 0$, we prove that a range exists.

In summary, no matter whether $V < vq_l - \frac{(1-\delta)Nn^2v^2q_l^2}{2\alpha_p}$ or not, a subscription fee range always exists which supports the separating equilibrium. \square

4.2.2. Pooling outcome: if all SPs get certified

If all SPs subscribe to certification services, they will all invest in regulative practices to control outbound traffic and they do not need to deploy protective practices to filter

inbound traffic from other certified SPs. Lemma 4 gives certified SPs' optimal strategies in the pooling outcome.⁵

Lemma 4. In the pooling outcome where all SPs are certified, the optimal strategies of a certified SP of type q_i are as follows:

- (1) it only invests in regulative practices and the level of effectiveness is $x_{ir}^{cp} = \frac{1}{\alpha_c} Mvq_i$;⁶
- (2) it charges $2VNn$ to its customers for Internet access services;
- (3) its profit is $\pi_i^{cp} = 2MV - Mvq_i + \frac{1}{2\alpha_r} (Mvq_i)^2 - t$.

Proof. In a pooling equilibrium where all SPs get certified, a customer is fully insured to enjoy two-way communication with all the other customers. Therefore, the SP will set price as $2VNn$. Including security investment, a SP with type q_i gains profit $\pi_i^{cp} = 2MV - Mv(1 - x_{ir}^{cp})q_i - C_r(x_{ir}^{cp}) - t$. Taking first order derivative over x_{ir}^{cp} yields optimal effectiveness $\frac{1}{\alpha_c} Mvq_i$. With x_{ir}^{cp} , we can obtain π_i^{cp} . \square

In this case, the size of the non-certified network diminishes to dimension 0, leaving only the certified network. All SPs take the responsibility for security and focus on the relatively more effective regulative practices. They are indifferent with whether to block or filter the traffic which is sent from the non-certified network because adopting either strategy yields the same expected payoff in equilibrium. In this paper, we assume that they block to induce stronger subscription incentives. Lemma 5 predicts the optimal strategies for a SP who deviates from equilibrium and stays non-certified.

Lemma 5. If a SP deviates from the pooling equilibrium and stays non-certified, it will invest in protective practices and the effective is $x_i^{np} = \frac{1}{\alpha_p} n^2 v q_l$ and its profit is

$$\pi_i^{np} = (M + n^2)V - MvE[q] + \frac{1}{\alpha_r} \left[(1 - \delta)(Mvq_l)^2 + \delta(Mvq_h)^2 \right] + \left[\frac{1}{2\alpha_p} - \frac{N}{\alpha_r} \right] n^4 v^2 q_l^2$$

Proof. The proof can be directly calculated based on the proof of Lemma 4. \square

When N is large, $M \gg n^2$. We therefore can simplify π_i^{np} using an approximate form $\pi_i^{np} = MV - MvE[q] + \frac{1}{\alpha_c} [(1 - \delta)(Mvq_l)^2 + \delta(Mvq_h)^2]$. In the pooling equilibrium, both high-type and low-type SPs should find it optimal to acquire a certification, compared to one's expected profit when it deviates. The subscription fee needs to be low enough to support such an incentive, as described in the following Proposition 4.

Proposition 4. The pooling equilibrium holds when

$$t \leq t^p = MV - Mv\delta(q_l - q_h) - \frac{1}{\alpha_r} \left[\delta(Mvq_h)^2 + \left(\frac{1}{2} - \delta \right) (Mvq_l)^2 \right]$$

Proof. In the pooling equilibrium, all the SPs should find it optimal to subscribe to the certification program. Since π_i^{cp} decreases as q_i increases, and π_i^{np} stays the same for both

⁵ Similar to the notation used in the separating outcome, we use the superscripts cp and np to denote the strategies and payoffs in the pooling outcome for certified and non-certified SPs respectively.

⁶ We use superscript p to represent the pooling outcome.

types, we only need to compare π_i^{CP} and π_i^{PP} when $q_i = q_l$, which yields the upper bound for the subscription fee t . \square

4.2.3. The efficiency of the certification mechanism

The role of the certification provider is to induce SPs to join in the certified network where each is responsible for the malicious traffic generated by its own users. The above section provides different ranges of subscription fees that support the separating and pooling equilibria. The certification provider can be a non-profit organization whose best interest is to induce the more effective security practices. Or, the certification provider can be self-interested and set a fee to maximize its total revenue. In this section, we will analyze the efficiency level of different equilibria. We then in Section 4.2.4 analyze the profitability of the certification provider.

Since the regulative practices are more efficient than the protective practices, we can conclude that a pooling outcome where both high-type and low-type SPs subscribe to the certification program would be the most efficient outcome. Now we define the efficiency level (E) as the total profit (excluding the subscription fee) gained by all the SPs. Then E in different cases can be calculated as follows:

$$E^b = N\pi^b = 2NMV - NMvE[q] + \frac{N}{2\alpha_p}(MvE[q])^2 \quad (\text{Benchmarkcase})$$

$$\begin{aligned} E^s &= \delta N(\pi_h^{CS} + t^s) + (1 - \delta)N\pi_l^{PS} \\ &= 2MNV - MNvE[q] + \frac{N}{2\alpha_p}(Mv(1 - \delta)q_l)^2 + \frac{N}{\alpha_r}\left[1 - \frac{\delta}{2}\right](Mv\delta q_h)^2 \\ &\quad - \delta(1 - \delta)NK^s\left[2MV - Mvq_l + \frac{(1 - \delta)}{2\alpha_p}(Mvq_l)^2\right] \end{aligned} \quad (\text{Separatingoutcome})$$

$$\begin{aligned} E^p &= \delta N\pi_h^{CP} + (1 - \delta)N\pi_l^{CP} + Nt^p \\ &= 2MNV - MNvE[q] + \frac{N}{2\alpha_r}\left[\delta(Mvq_h)^2 + (1 - \delta)(Mvq_l)^2\right] \end{aligned} \quad (\text{Poolingoutcome})$$

Proposition 5. Comparing E^b, E^s , and E^p , we have $E^p > E^b$, $E^p > E^s$, and $\lim_{\delta \rightarrow 1} E^s = E^p$.

Proof. The result can be derived from direct comparisons of the formulae of E^b, E^s , and E^p . \square

The pooling outcome is always better than the other two outcomes. The separating equilibrium can be relatively efficient when δ is large enough, that is, when most of the SPs are high-type SPs and the size of the non-certified network is relatively small. In today’s networking environment, this is generally true.

4.2.4. Profitability of certification provider

The remaining question is “who should assume the role of a certification provider?” Or more importantly, “what is the certification provider’s objective?” From a social efficiency point of view, the certification provider should induce the most efficient outcome (i.e. the pooling outcome) to create a safe Internet. However, if the certification provider is concerned about its own profit (i.e. the total revenue collected by the subscription fee), then it may find it more profitable to

induce the separating outcome. We will analyze the certification provider’s choice of the subscription fee to maximize its profit and consequent equilibria in this section.

If the certification provider attempts to maximize profit, it will only set a positive subscription fee. Lemma 6 shows whether the upper bounds of the ranges described in Propositions 3 and 4 are positive.

Lemma 6. t^p and t_{k2}^s are both positive, and t_{p2}^s is negative.

Proof.

- (1) We first show that t^p is positive: substituting $\delta=0$ and $\delta=1$ to t^p respectively, we know that $t^p(\delta=0) > 0$ and $t^p(\delta=1) > 0$. In addition, $\frac{dt^p}{d\delta} = \left[\frac{M^2v^2(q_l+q_h)}{\alpha_r} - 1\right][Mv(q_l - q_h)]$ is not a function of δ . We then conclude that $t^p > 0$ for all $\delta \in [0,1]$.
- (2) Next, we show that $t_{k2}^s > 0$ (when $V < vq_l - \frac{(1-\delta)Mn^2v^2q_l^2}{2\alpha_p}$): since we have $MV > \frac{1}{2}\alpha_r > \frac{1}{2\alpha_r}(Mvq_h)^2$ and $V < vq_l - \frac{(1-\delta)Mn^2v^2q_l^2}{2\alpha_p}$, we can rearrange the terms of t_{k2}^s ,

$$\begin{aligned} t_{k2}^s &= MV(2\delta - 1) + Mv(1 - \delta)q_l - \frac{(Mv)^2}{2}\left[\frac{(\delta q_h)^2}{\alpha_r} + \frac{((1 - \delta)q_l)^2}{\alpha_p}\right] \\ &= MV\delta - \frac{1}{2\alpha_r}(Mv\delta q_h)^2 + (1 - \delta)\left[Mvq_l - MV - \frac{(1 - \delta)(Mvq_l)^2}{2\alpha_p}\right] > 0. \end{aligned}$$

- (3) t_{p2}^s is negative, which is straight-forward. \square

Based on the result of Lemma 6, we conclude that only a pooling equilibrium will be induced by the certification provider if $V \geq vq_l - \frac{(1-\delta)Mn^2v^2q_l^2}{2\alpha_p}$. The maximal price that can be charged is t^p and the total profit is $t^p \cdot N$. If $V < vq_l - \frac{(1-\delta)Mn^2v^2q_l^2}{2\alpha_p}$, the certification provider may choose to charge t^p (profit $t^p \cdot N$) or t_{k2}^s (profit $t_{k2}^s \cdot \delta N$) depending which is more profitable.

Proposition 6. When $V < vq_l - \frac{(1-\delta)Mn^2v^2q_l^2}{2\alpha_p}$, there exist two critical values, δ_1 and δ_2 such that $0 < \delta_1 < \delta_2 < 1$, and:

- 1) The certification provider sets $t = t^p$ and all the SPs will subscribe if $\delta \in [0, \delta_1)$;
- 2) The certification provider sets $t = t_{k2}^s$ and only high-type SPs will subscribe if $\delta \in [\delta_2, 1)$.

Proof. Define $\Pi^p = t^p N$ and $\Pi^s = t_{k2}^s \delta N$ as the certification provider’s profit. Moreover, define $\Delta = \Pi^p - \Pi^s$. It can be shown that $\lim_{\delta \rightarrow -\infty} \Delta < 0$, $\Delta(\delta=0) > 0$, $\Delta(\delta=1) < 0$ and $\lim_{\delta \rightarrow +\infty} \Delta > 0$. In addition, Δ is continuous in δ and the highest order is δ^3 . Hence, there exists only one $\delta \in (0,1)$ such that $\Delta(\delta) = 0$. If $\delta \in [0, \delta)$, then $\Delta(\delta) > 0$ and the pooling outcome is more profitable. If $\delta \in (\delta, 1]$ then $\Delta(\delta) < 0$ and the separating outcome is preferred.

In order to successfully induce the pooling (separating) outcome, we also need to have $t^p > t_{k2}^s$ ($t^p < t_{k2}^s$). Define $\Delta' = t^p - t_{k2}^s$. We have that $\Delta'(\delta=0) > 0$, $\Delta'(\delta=1) < 0$, and Δ' is strictly convex. Hence there exists a $\delta \in (0,1)$ such that $\Delta'(\delta) = 0$. We can also conclude that $\delta < \delta$ since $\Delta(\delta) = t^p(\delta)N - t_{k2}^s(\delta) \cdot \delta N = N[\Delta'(\delta) + (1 - \delta) \cdot t_{k2}^s(\delta)] = N \cdot (1 - \delta) \cdot t_{k2}^s(\delta) > 0$. Define $\delta_1 = \delta$ and $\delta_2 = \delta$, we can draw the conclusion of Proposition 6. When $\delta \in [\delta, \delta]$, the certification provider may not be able to induce the more profitable pooling equilibrium since a separating equilibrium can also exist in that subscription fee range. \square

Proposition 6 shows that the certification provider will induce the separating outcome only when the number of high-type SPs is large enough. Proposition 5 shows that the efficiency level of the separating outcome is high if δ is large. Therefore, we believe a profit-maximizing certification provider will still induce some level of efficiency in our proposed structure.

5. Discussion and conclusion

This research examines the Internet architecture and address Internet security issues from an economic perspective. We propose a certification mechanism to induce SPs to exert collective efforts and improve Internet security. To be more specific, the proposed mechanism provides certified SPs incentives to deploy regulative practices. We use a game-theoretic model to examine the efficiency of our mechanism. The results show that our mechanism can increase the efficiency for all the Internet Service Providers. By providing SPs with appropriate incentives, our mechanism can create a better communication environment over the Internet.

The challenging issue is, *who should be the certification provider?* The certification provider can be a non-profit institution, such as Internet Corporation for Assigned Names and Numbers (ICANN). ICANN is a central authority with limited power in the essentially decentralized and neutral global network. However, its functions are restricted to running the addressing system, giving out blocks of unique identifiers to countries and private registries. Commentators have suggested ICANN should play an enhanced role in governing the unregulated Internet. By providing certification services, it introduces a soft regulation to the Internet, characterized by the fact that participation is voluntary, and participants choose their actions based on self-interest. The certification provider can also be a for-profit organization. The previous analysis examines the certification provider's profit and discusses its impact on the overall efficiency level.

One concern of our proposed mechanism is that if the certified network completely blocks inbound traffic sent from the non-certified network, then the overall network will suffer. Our result shows that blocking is an optimal strategy in the separating outcome only when the value of communication V is relatively low compared to the disutility caused by the malicious attack. In the pooling outcome, everyone will join the certified network and "blocking" is only a threat to those who deviate. Alternatively, we suggest the certified SPs consider strategies such as slowing down the incoming traffic sent from non-certified networks to deteriorate the non-certified SPs' payoff. However, due to the network interdependency, the certified SPs will also suffer from such a strategy. How to provide the certified SPs proper incentives to "punish" those non-certified ones deserves further study.

This paper characterizes the effectiveness of security practices using a single parameter, x , representing the *false negative*⁷. In most control settings, both false negatives and false positives are used to describe the effectiveness of security

practices. Since regulative practices generally outperform protective practices in reducing errors, the analysis and results considering both false positive and false negative will be similar.

The implementation of the certification mechanism may generate extra overhead to identify the service providers' certification status. We ignored such an impact in our model by assuming that the size of overhead is negligible compared to the regular traffic. In situations that the assumption does not hold, we suggest the certification provider to adjust the subscription fee to accommodate the overhead cost. Although the overhead will create a deadweight loss which reduces the value of the certification mechanism, the loss is inevitable as no security mechanism is free. Given the rising concerns on security, the overhead should not stop the implementation of the certification mechanism.

The main contribution of our paper is to propose a new incentive framework to the management of network security. Compared to the current Internet infrastructure which is open to everyone and consequently leaves everyone exposed to the security risks, our vision of the Internet is one where all the active parties (e.g. the service providers) should work collectively as a whole to detect potential security risks and eliminate the possible damage at the earliest stage. Our proposed framework also suggests possible exclusion of the incompetent service providers who cannot afford to make such an endeavor under certain conditions. Such a proposal may sound controversial from an idealistic point of view. However, it can induce those competent service providers to take more active actions in safeguarding the Internet, providing the individual users a worry-free environment. Our analytical results prove that the framework produces more efficiency for Internet communication.

Acknowledgements

We thank Dr. Manoj Parameswaran from the Department of Operations & Management Information Systems in Santa Clara University for his insightful comments, Dr. Ashish Arora from Carnegie Mellon University for suggesting us to clarify the novelty of our game-theoretic model, Dr. Ruhul Telang from Carnegie Mellon University for providing us helpful references. We also thank Dr. Vitaly Shmatikov, Dr. Benjamin J. Kuipers, Dr. Mohamed G. Gouda, and Dr. Lili Qiu from the Department of Computer Sciences at the University of Texas at Austin, Dr. Claire Vishik from Intel, John S. Quarterman from InternetPerils, Inc., as well as seminar participants at the University of Texas at Austin, the Fourth Workshop on eBusiness at Las Vegas, the Sixteen Workshop of Information Technologies and Systems at Milwaukee. We are responsible for all possible errors.

References

- [1] R. Anderson, T. Moore, The economics of information security, *Science* 314 (October 2006) 610–613.
- [2] T. August, T. Tunca, Network software security and user incentive, *Management Science* 52 (11) (2006) 1703–1720.
- [3] H. Ballani, P. Francis, X. Zhang, A study of prefix hijacking and interception of the Internet, *Proceedings of ACM SIGCOMM 2007*, August 2007.
- [4] H. Cavusoglu, B. Mishra, S. Raghunathan, A model for evaluation IT security investments, *Communications of the ACM* 47 (7) (2004) 87–92.

⁷ A false negative is the polar opposite of a false positive. A false negative occurs when the security protection fails to detect a malicious activity. For example, a virus scanner fails to detect a virus in an infected file or an email filter fails to detect a spam email.

- [5] H. Chan, D. Dash, A. Perrig, H. Zhang, Modeling adaptability of secure BGP protocols, Proceedings of ACM SIGCOMM 2006, September 2006.
- [6] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, A. Rubin, Working around BGP: an incremental approach to improving security and accuracy in interdomain routing, Proceedings of Symposium on Network and Distributed System Security (NDSS'03), February 2003.
- [7] L.A. Gordon, M.P. Loeb, The economics of information security investment, ACM Transactions on Information and System Security 5 (4) (2002) 438–457.
- [8] Y.-C. Hu, A. Perrig, M. Sirbu, SPV: Secure patch vector routing for securing BGP, Proceedings of ACM SIGCOMM 2004, September 2004.
- [9] C.D. Huang, Q. Hu, R. Behara, Economics of information security investment in the case of simultaneous attacks, Proceedings of the Workshop on the Economics of Information Security (WEIS2006), (Cambridge, UK, 2006).
- [10] Y. Huang, X. Geng, A.B. Whinston, Defeating DDoS attacks by fixing the incentive chain, ACM Transactions on Internet Technology 7 (2) (2007).
- [11] K. Kannan, R. Telang, Market for software vulnerabilities? Think again", Management Science 51 (5) (2005) 726–740.
- [12] S. Kent, C. Lynn, J. Mikkelsen, K. Seo, Secure border gateway protocol (S-BGP) — real world performance and deployment issues. Proceedings of the Network and distributed Systems Security Symposium (NDSS 2000), February 2000, pp. 103–116, San Diego, CA.
- [13] S. Kent, C. Lynn, K. Seo, Security border gateway protocol (S-BGP), IEEE Journal of Selected Areas in communications 18 (4) (April 2000) 582–592.
- [14] H. Kunreuther, G. Heal, Interdependent security, Journal of Risk and Uncertainty 26 (2/3) (2003) 231–249.
- [15] D. Lichtman, E. Posner, Holding Internet Service Providers accountable, John M. Olin Law & Economics Working Paper No. 217, http://ssrn.com/abstract_id=573502, (2004).
- [16] H. Ogut, N. Menon, S. Raghunathan, Cyber insurance and IT security investment: impact of interdependent risk, Proceedings of the Workshop on the Economics of Information Security (WEIS2005), Harvard University, Cambridge, MA, 2005.
- [17] M. Parameswaran, X. Zhao, A.B. Whinston, F. Fang, Reengineering the Internet for better security, IEEE Computer 40 (1) (2007) 40–44 (January 2007).
- [18] H.R. Varian, H.R. Managing online security risks, NY Times, <http://www.ischool.berkeley.edu/~hal/people/hal/NYTimes/2000-06-01.html>, (2000).
- [19] R. White, Securing BGP through secure origin BGP, Technical Report, Cisco Internet Protocol Journal, September 2003.

Xia Zhao is a research fellow at the Tuck School of Business, Dartmouth College. Her research interests include electronic commerce, Internet security, and electronic communities. Zhao received her PhD in Information Systems from the University of Texas at Austin. Contact her at xia.zhao@dartmouth.edu.

Fang Fang is an assistant professor in the College of Business Administration at California State University San Marcos. Her research interests include knowledge markets, financial markets, and network finance. Fang received her PhD in Information Systems from the University of Texas at Austin. Contact her at fangfang@csusm.edu.

Andrew B. Whinston is the Hugh Roy Cullen Centennial Chair Professor in Information Systems at the Graduate School of Business in the University of Texas at Austin, where he also teaches economics and computer science and directs the Center for Research in Electronic Commerce. His research interests include electronic commerce, knowledge management, online auctions, and financial markets. Whinston received a PhD in management from Carnegie Mellon University. Contact him at abw@uts.cc.utexas.edu.