Decision Support Systems 51 (2011) 745–759

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

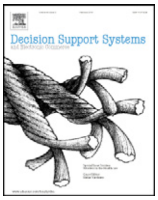

Predicting stock market returns from malicious attacks: A comparative analysis of vector autoregression and time-delayed neural networks

Lara Khansa [a,*], Divakaran Liginlal [b]

[a] Department of Business Information Technology, Virginia Polytechnic Institute and State University, Blacksburg, VA, United States
[b] Carnegie Mellon University in Qatar, Doha, Qatar

ARTICLE INFO

Available online 1 February 2011

Keywords:
Malicious attacks
Stock price
Time-delayed artificial neural network
Vector autoregression

ABSTRACT

With the growing importance of Internet based businesses, malicious code attacks on information technology infrastructures have been on the rise. Prior studies have indicated that these malicious attacks are associated with detrimental economic effects on the attacked firms. On the other hand, we conjecture that more intense malicious attacks boost the stock price of information security firms. Furthermore, we use artificial neural networks and vector autoregression analyses as complementary methods to study the relationship between the stock market returns of information security firms and the intensity of malicious attacks, computed as the product of the number of malicious attacks and their severity levels. A major contribution of this work is the resulting time delayed artificial neural network model that allows stock return predictions and is particularly useful as an investment decision support system for hedge funds and other investors, whose portfolios are at risk of losing market value during malicious attacks.

1. Introduction

Malicious code consists of software or firmware intended to execute an unauthorized computer process with the purpose of disrupting the target's information systems [48]. Viruses, worms, Trojans, and other code based entities that infect a host fall into this category. The growth of Internet based businesses and globalization have exposed information technology (IT) infrastructures of firms to malicious code attacks of increasing intensity. In response, information security firms, i.e. the vendors of products that protect information systems from such malicious attacks, have invested heavily in research and development to come up with more resilient products. While malicious attacks have been associated with detrimental economic effects on attacked firms [8,10,22,23,30,40], we conjecture that they are beneficial to the market value of information security firms. The basis for this argument is that higher intensity of malicious attacks leads to higher investments in information security products and services [38], which, in turn, should have a positive effect on the stock price of the information security firms who sell these products and services. In this paper, we attempt to investigate two research questions: (i) *Are malicious attacks beneficial to the stock price of information security firms*, and *(ii) How well can the stock market returns of information security firms be predicted from these malicious attacks?*

Artificial neural networks (ANNs) have been widely used in the areas of intrusion detection and prevention, spam control, and financial prediction. The unpredictability in the arrival and intensity of malicious attacks makes it ideal to exploit the adaptability of ANNs to forecast the stock market returns of information security firms, given a certain level of intensity of malicious attacks. Our literature survey indicated that prior studies have not attempted to relate the market performance of information security firms to the intensity levels of malicious attacks, computed as the product of the severity of malicious attacks, which are assessed by antivirus providers' subject matter experts, and the number of malicious attacks. By adopting a complementary approach using both vector autoregression (VAR) methods and ANNs, we attempt to fill this gap in the literature.

Forecasting models proposed in the information security economics literature include risk management models [38,54], game theoretical models [11], real options (RO) based models [37], and event study analyses [10,22]. Risk management models generally take into account both a risk analysis phase and a risk assessment phase. The risk analysis phase involves the identification of threats to system security, the probability of their occurrence, and their resulting impact. The risk assessment phase consists of identifying safeguards to mitigate the impact of the threats, followed by a cost benefit analysis. Commonly used risk metrics such as Average Loss Expectancy (ALE) have been criticized for being overly complex when they attempt to address all threats and vulnerabilities [54]. RO based models that quantify the benefits of information security investments

[*] Corresponding author at: Pamplin College of Business, Department of Business Information technology, 2062 Pamplin Hall (0235), Blacksburg, VA 24061–0101, United States. Tel.: +1 540 231 5003.
E-mail addresses: larak@vt.edu (L. Khansa), liginlal@cmu.edu (D. Liginlal).

0167-9236/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.dss.2011.01.010

have also been developed [37]. However, these RO based models impose assumptions on the distribution of malicious attacks. In contrast, our analyses in this paper are based on actual records of malicious attacks and hard financial stock market data. Further, event study analyses have been conducted to relate security breaches resulting from malicious attacks to the stock price of information security firms [10,22]. Event studies only capture the momentary impacts of security breaches around the time of their announcements. In this paper, we relate the time series of malicious attacks to that of the stock market returns of information security firms and study their dynamic relationship over time. We first make use of VAR analysis to investigate this relationship, in the Granger causality sense [24]. VAR analysis, in the context of this paper, is a better fit than other regression methods because it captures time lagged effects and feedback between variables. These effects and feedback are particularly relevant in time series data where relationships between variables could go both ways. One could argue that when the market value of information security firms increases, their ability to innovate and reduce the intensity of future malicious attacks increases as well. Sims [52] advocated the use of VAR models for such analyses because they make no a priori assumptions regarding the relationships between variables, thus avoiding the identification restrictions of structural models. After we have established the significance between our model's variables and the time lag using VAR analysis, we use ANN based analytical methods that complement VAR analysis by effectively mapping complex nonlinearities without specification of assumptions regarding the statistical distribution or properties of the underlying data.

The remainder of this paper is organized as follows. Section 2 provides a background to the research in the subject disciplines, surveys contemporary literature, and develops our hypothesis. Section 3 covers the research design, including a discussion of the data collection methods and measures. VAR analysis and results are presented in Section 4. The time delayed ANN based analysis, comparisons with the results of VAR, and sensitivity analyses with respect to ANN parameters and data quality are presented in Section 5. Section 6 discusses the implications, contributions, and limitations of the study, and avenues for future research.

2. Theoretical background

Malicious attackers can be classified as one of three types: a masquerader who is not a legitimate user but attempts to exploit a legitimate user account, a misfeasor who is a legitimate user but attempts to access resources to which he/she does not have usage permission, and a clandestine user who attempts to gain administrative control of the system [56]. What all these types of intruders have in common is that they all need a way into the system they desire to access. There exist many ways in which to gain access to a system, such as tricking a legitimate user over the telephone or by email into releasing account details. However, one of the most popular means of gaining entry is by exploiting software bugs, also called vulnerabilities, in host operating systems or Web applications that are running on Web servers. Software vendors have been shown to follow the strategy of releasing their products early and fixing them later through patching [3,44], especially when the market size and the degree of competition are high. Spier [55] showed that if manufacturers have the ability to repurchase their defective product, they have fewer incentives to design safer products. In a sense, software vendors have the ability to “repurchase their product” or more accurately they can repair their product without a physical recall. They can “ship” a patch to a user and have the user “install” the solution. The fact that software vendors do not have to physically recall their products from the customer could lessen their incentives to reduce vulnerabilities. Anderson [2] concluded that software vendors are capable of creating more secure software but that the economics of the software industry provide few incentives to encourage the development of more secure products and that the idea of shipping the product now and fixing the bugs later is a perfectly rational approach. The number of software vulnerabilities has increased dramatically over the years. Even in recent years, Symantec, a major antivirus provider (http://www.symantec.com), reported a 19% increase in documented vulnerabilities from 2007 to 2008, and that, in parallel, new malicious code signatures have increased by an astonishing 265% in 2008 compared to 2007.

A parallel exists between vulnerabilities in software and vulnerabilities in manufactured products. This paper, therefore, draws upon the research stream studying the effect of product defects on the market value of related firms, whether attacked firms or firms producing solutions for these defects. Bad reputation and lack of customer satisfaction have been shown to affect future customer behavior and, in turn, the level, timing, and risk of future cash flows of suppliers [5,12,29,32,33,35,48,50,51,63]. Chen et al. [12] showed that firms are better off being passive in responding to product recalls because proactive strategies are perceived by investors as a signal of larger financial losses to the firm. Rhee and Haunschild [48] tested automobile recalls from 1975 to 1999 and found that highly reputable firms are punished more upon product recalls than less reputable firms. Rupp [51] found that recalls initiated by the government were as damaging to shareholders as other recalls. White and Pomponi [63] reported the costs of recalls to consumer products companies to be more than 6 billion US$ per year. Kamp and Burton [35] reported that Medtronic's fourth quarter 2008 earnings fell 69% due to charges of product flaws and a safety notice. Mattel recalled millions of toys in 2007 and reported a cost of 30 million US$ related to those recalls. The costs (tangible and intangible) to Toyota of their recent recalls are expected to be astronomical. Early estimates of the cost of these recalls reach 2 billion US$ [32]. There are several studies that illustrate the negative effects of product recalls on firm value. Jarrell and Peltzman [33] found that drug recalls by the Food and Drug Administration (FDA) resulted in an average 6% loss in stock equity values for the affected firms. Moreover, some of the effects of the recall on stock equity values spilled over to other drug companies not directly affected by the recall. Similarly, Rubin et al. [50] estimated that product recalls by the U.S. Consumer Product Safety Commission resulted in an average 7% reduction in the stock equity values of the firms involved. Hendricks and Singhal [29] recounted the negative effects of supply chain glitches that resulted in production and shipment delays, on the market value of suppliers.

In the particular realm of information security, several researchers have investigated the impact of information security breaches on the market value of affected firms. Goel and Shawky [23] studied security breaches over the period 2004 2008 and noted a significantly negative impact on the market value of breached firms. Campbell et al. [8] found that only breaches linked to loss of confidential information had significant negative effects on the stock price of firms, while the impact of non confidential breaches were not significantly different from zero. Hovav and D'Arcy [30] also studied the impact of denial of service attack announcements on the market over 4.5 years and showed that the market penalizes “Internet specific” companies more than other companies. Liginlal et al. [40] found that investors' confidence in a financial firm's continuity is particularly abated after a human error related privacy breach. Cavusoglu et al. [10] studied the change in market value of firms whose systems had been breached. The study showed that the announcement of a security breach decreased the market capitalization values of a firm, on average, by 2.1 billion US$ within two days of the breach. Furthermore, the study demonstrated a significant information transfer effect to information security firms. Garg et al. [22] reported similar findings and confirmed this transfer effect. Both Cavusoglu et al. [10] and Garg et al. [22] used event study analyses to uncover the momentary positive transfer effect on the stock price of information security firms. An event study analysis quantifies the effect of an event, in this case the occurrence of a malicious attack, over a short time interval around the day when the attack occurred. It does not, however, analyze the longer term effects of these attacks. Our research questions go beyond the short time span around the attack to suggest that the effect of malicious attacks on the stock price of information security firms is of longer term.

Clemons [15] argued that some IT investments are made to limit future losses rather than to generate value. Investing in information security fits this characteristic perfectly. General deterrence theory in criminology suggests that criminal behavior can be influenced by deterrence certainty, the likelihood of being punished, and deterrence severity, the perceived harshness of the punishment [6,18]. The results of Khansa and Liginlal [38], which showed that investment in information security is effective in reducing the intensity of malicious attacks, also support deterrence theory. Cavusoglu et al. [11] showed that firms can benefit the most from their information security investments when they move preemptively ahead of hackers and are able to learn from prior malicious attacks. In their 2004 paper, Cavusoglu et al. [10] reported that, "…as firms invest more on security, demand for security products goes up. Gartner predicts that U.S. companies' investments in information security will increase from the current 0.4% of revenue to 4% of revenue by 2011, a 1000% increase [49]." Similarly, it has been shown that higher expected revenues for information security firms raise investors' beliefs about the sustainability of the information security market sector [38]. In other words, firms invest in information security to counteract the potential nefarious effects of expected future attacks. Their investments in information security help boost the future expected revenues and profits of information security firms, and are, consequently, positively reflected on information security firms' stock prices. This not only indicates that the market performance of the information security sector is demand driven, but also confirms the results in [21], which showed that stock prices have been more sensitive to demand driven output fluctuations than to supply driven variations. Our hypothesis is based on the semi strong version of the efficient market hypothesis (EMH) [19,20], which states that as new publicly available information is received it is gradually absorbed by investors and incorporated into share prices. Although one might think that actual investments in information security products and services take a long time to be reflected on the revenues of information security firms, stock market investors' reactions are swift because they are based on the expected future increase in the revenues of information security firms, rather than on the realized increase in these revenues. Basically, the potential for increase in revenues is factored in the stock prices upon announcement of the malicious attacks. In fact, prior research [10,22] showed this quick response based on event study analyses. As such, we hypothesize the following: *A higher intensity of malicious attacks is associated with an increase in the stock market returns of information security firms*, as represented by the dashed line in Fig. 1.

Numerous studies have proposed ANNs as a viable and even better alternative than regression analysis for classification and prediction problems. Some studies compared neural network with regression analysis and showed how "neural network analysis represents an exciting and complementary computational methodology for cross cultural researchers ([60], p. 232)." Zhou et al. [67] compared the effectiveness of four strategies for automatically detecting deceptive computer communication and suggested a neural network based approach as particularly promising. Sinha and May [53] compared the performance of neural networks with four other data mining techniques, including logistic regression, and found that neural networks, together with logistic regression, showed better generalizability, scalability, and robustness than linear discriminant analysis, decision trees, or nearest neighbor techniques. Chiang et al. [13] showed that neural networks are significantly better than logistic regression models in terms of predicting power. Similarly, Boyacioğlu et al. [7] examined neural networks, support vector machines, and multivariate statistical methods to predict bank financial failures. They found that multilayer perceptron and learning vector models are good for forecasting. Alfaro et al. [1] presented AdaBoost, an improved neural network classification algorithm to predict bankruptcy in a set of European firms. Similarly, Walczak and Sincich [61] compared results of their neural network model to that of logistic regression analysis for modeling student enrollment decision making at a small private university. The study showed that much improvement was realized via neural networks, as neural networks effectively halved the student applicant load for each counselor. Lu et al. [41] used a combination of support vector regression (SVR), based upon the support vector machine (SVM) neural network algorithm, and independent component analysis (ICA) to forecast financial time series. For a comprehensive comparison between neural networks and a multitude of other statistical techniques, please refer to [45,62]. The authors conclude that artificial neural networks have immense

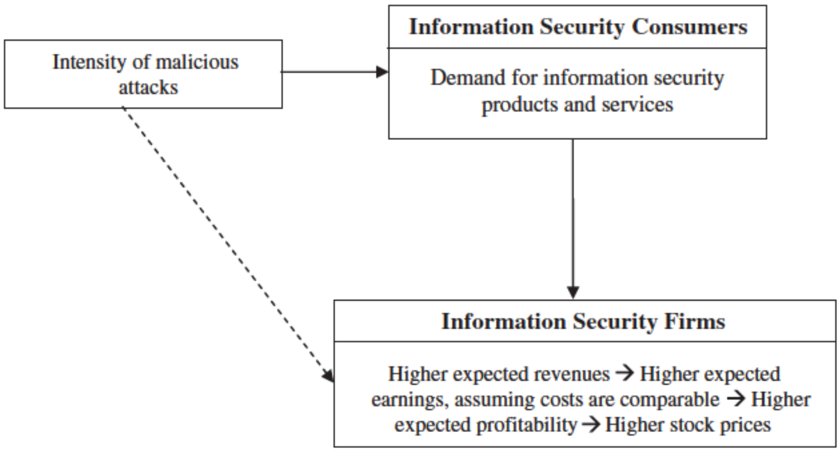
Intensity of malicious attacks
Information Security Consumers
Demand for information security products and services
Information Security Firms
Higher expected revenues → Higher expected earnings, assuming costs are comparable → Higher expected profitability → Higher stock prices

Fig. 1. Malicious attacks and stock price of information security firms.

potential as a tool for classification and prediction because they can automatically approximate any nonlinear mathematical function, which is particularly useful when the relationship between variables is unknown or complex. None of the studies surveyed in [45,62] used VAR analysis that is particularly efficient at modeling time lagged effects. In this paper, we conduct both VAR and ANN based analyses as complementary methods that together strengthen the veracity of our results.

3. Research design

3.1. Data collection: sample of information security firms

We selected all public information security firms listed on Yahoo! Finance (www.finance.yahoo.com). We chose the time period of 1996 to 2008 for data collection because it encompasses a diverse range of economic conditions for technology related firms. The 1996 2008 period is diverse in the sense that it includes the pre bubble years, the years of the Internet bubble, the quasi recession of 2001, the subsequent short economic recovery, as well as the credit crisis of the late 2007 and 2008. The inclusion of all these economic settings adds to the generalizability of our results. To start with, we analyzed firms in the security software and services industry sector. Firms were then added from other industry sectors, namely the Internet software and services sector that includes firms such as Symantec and VeriSign; the computer peripherals sector that includes companies such as Secure Computing; and the networking and communication devices sector that includes firms such as Blue Coat Systems. We made sure that information security firms cited in the list of nominees of the 2009 Information Security Excellence Awards by Information Security Magazine were all accounted for in our sample. We also added public information security firms from the list of attendees (including exhibitors and sponsors) at the 2009 RSA conference (http://www.rsaconference.com/), a global information security conference that attracts security firms and security professionals worldwide. We excluded from our sample firms with no SEC filings (SEC stands for Securities and Exchange Commission, a government agency responsible for the supervision and regulation of the securities industry), non U.S. based firms, and firms that are not listed in the Center for Research in Security Prices (CRSP) database, from which we collected all financial information (The CRSP database is maintained at the University of Chicago through the Wharton Research Data Services, WRDS (https://wrds.wharton.upenn.edu/). We ended up with a total of 88 public information security firms shown in Table A.1 in Appendix A. Table A.1 categorizes these information security firms into antivirus (AV) firms, network security firms, and identity and access management (IAM) firms. We computed the average of the daily stock prices of the firms in the resulting sample, at the close of each trading day, and obtained the time series in Fig. 2. The daily stock market returns were then computed and used in the VAR and ANN analyses.

3.2. Data collection: sample of malicious attacks

Micro data pertaining to the damaging extent of malicious attacks are not available directly from firms, as corporate decision makers are unwilling to fully disclose the dollar impact of malicious attacks on their IT infrastructures. We gathered around 11,000 malicious attacks from Symantec's website. We also collected each attack's severity impact that ranged from 1 to 3. The severity of malicious attacks is based on three criteria: (1) wildness level, (2) damage level, and (3) distribution level [38]. The wildness level is based on the number of infected machines, the number of infected sites, the lack of threat containment, and the difficulty of removal. Thurstone [58] justified using interval level measurements based on the law of comparative judgment. Similarly, the probabilistic Rasch model [47] provides a theoretical justification for obtaining interval level measurements from counts of observations such as assessments. Fig. 3 shows that most reported malicious attacks are Trojans and worms, especially after the IT bubble of the year 2000.

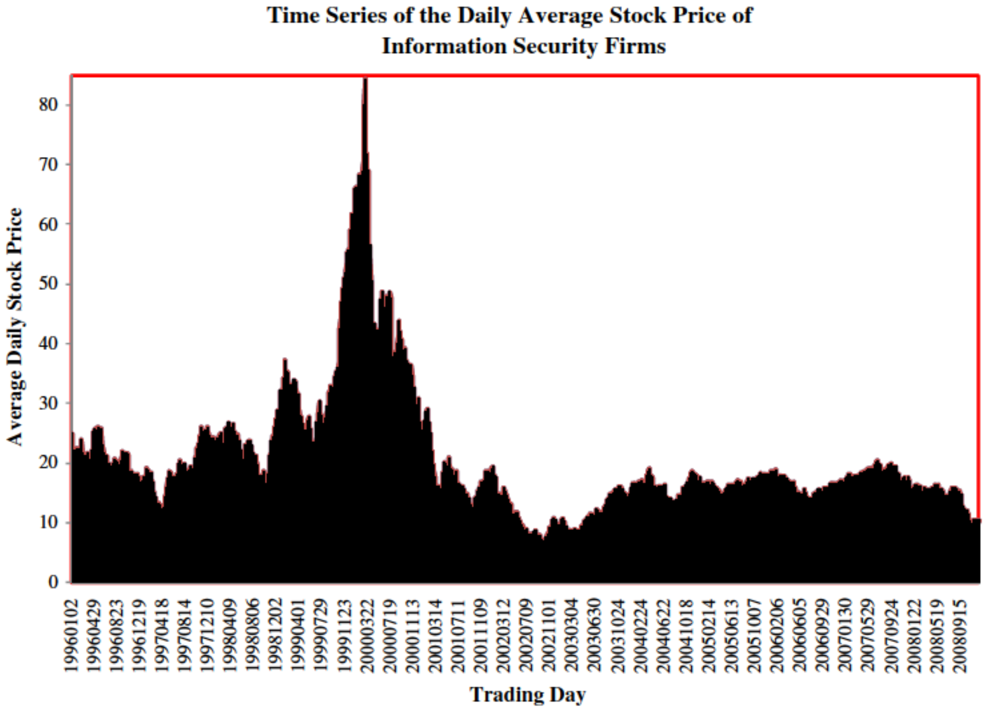

Fig. 2. Time series of the daily average stock price of information security firms.

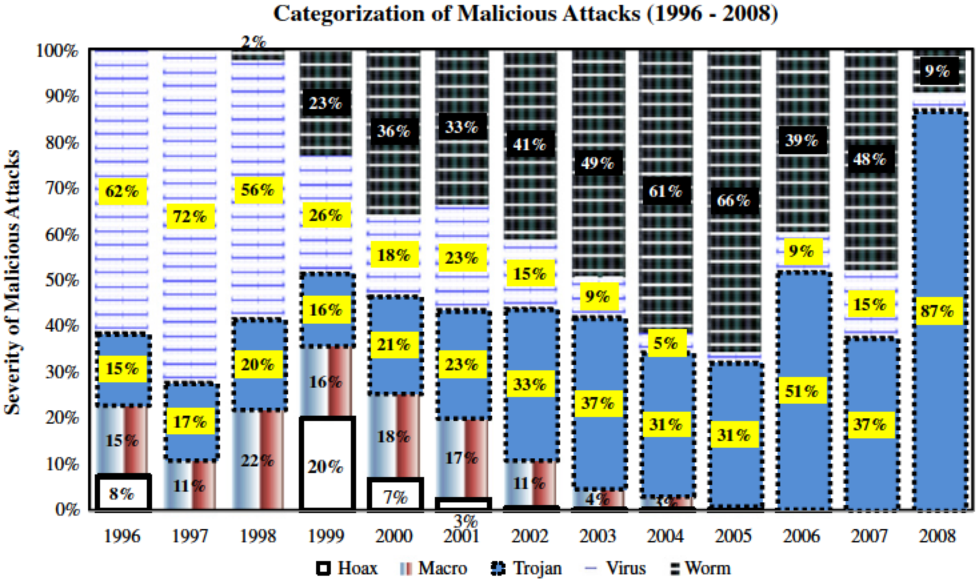

Fig. 3. Categorization of malicious attacks.

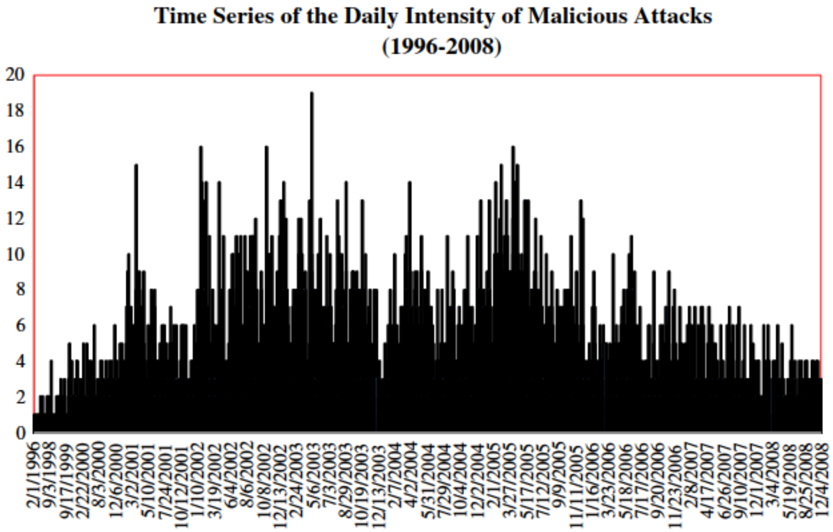

Fig. 4. Time series of the daily intensity of malicious attacks.

After collecting all malicious attacks and their impacts, we used the pivot table capability of Excel 2007 to aggregate all the intensity levels daily. For example, if 3 attacks occurred on a trading day with respective severity levels of 1, 2, and 3, the malicious intensity that occurred on this trading day is 6. Since there are only 252 trading days in a year, we filtered out the malicious attacks that happened to occur on a non trading day to obtain a malicious intensity time series that we can relate to the time series of stock market returns. We ended up with 7291 malicious attacks. The time series of the daily intensity of these malicious attacks is shown in Fig. 4.

Table 1 plots the distribution of the severity levels among the various attacks, revealing that only few attacks have a severity level of 3. The majority of attacks had severity levels of 1 or 2. Table B.1 in Appendix B gives the time series of the daily number of malicious attacks and the intensity of malicious attacks. We further correlated the two time series, i.e., one for malicious intensity and the other with only the number of malicious attacks, and found them to be highly correlated (0.9813).

By relating the intensity of malicious attacks, based on both the count of malicious attacks and the severity of malicious attacks as described above, we tailor the analyses to the viewpoint of information security firms, rather than the attacked firms. In fact, since we are interested in the effect of malicious attacks on the stock market returns

Table 1
Severity levels of malicious attacks.

Severity level	Number of malicious attacks
1	5703
2	1582
3	6

of information security firms, what really matters are the revenues that these firms accumulate as a result of malicious attacks.

3.3. Measures

The variables used in the models are summarized below.

1. The security return index. This is constructed by taking the average daily stock market returns of every information security firm in the sample. This security return index represents the market return of the information security sector as a whole.
2. The market index. This controls for prevailing market conditions including both macroeconomic and industry effects. The market index is constructed from the CRSP daily NYSE/AMEX/NASDAQ value weighted return. Prior research [39,40,42,43] recom mended using the CRSP market index to account for economic factors.
3. Malicious intensity. This is obtained by aggregating the intensity of malicious attacks daily over trading days from 1996 to 2008. To cross validate our malicious intensity data, instances of malicious attacks and the subjective estimates of their impacts from two other major antivirus vendors, namely McAfee and Trend Micro for the period 1996 to 2008, were collected. Upon aggregating the impact estimates daily, a high positive correlation was observed among the three datasets (highest of 0.897 between Symantec and McAfee), which suggests a consensus among the subjective assessments of all three sources and reinforces the credibility of the collected impact estimates.

Table 2 gives descriptive statistics of the three variables: the daily average stock market returns of information security firms, the daily intensity of malicious attacks, and the daily average market index.

There are 3275 input observations in all corresponding to the total number of days the market was actively trading from January 1996 to December 2008 and during which there occurred at least 1 malicious attack.

Table 2
Descriptive statistics.

Descriptive statistics	Daily average stock return of information security firms	Daily intensity of malicious attacks	Daily average stock market return
Mean	0.000144899	3.445753759	0.005292703
Standard error	0.000499167	0.053468769	0.00339691
Standard deviation	0.028557395	2.652503659	0.042291183
Minimum	0.203982877	1	0.186398
Maximum	0.164146406	19	0.09461

4. Vector autoregression analysis

A critical requirement before using VAR analysis is for all time series to be stationary and not to exhibit seasonality [46]. All three time series were analyzed for autocorrelation, trend, and seasonality before conducting further analyses. Initial plots of the partial autocorrelation function (PACFs) of the stock price time series indicated first order autocorrelations, reinforcing the discussion in the literature that there exists positive momentum in some stock price data, particularly at high frequencies. A similar pattern in malicious intensity data conjectures that the intensity of malicious attacks also follows a geometric random walk process, confirming the findings in [37]. However, when we derived the stock returns from the stock prices of information security firms, we found that the security return index time series is stationary. Similarly, the market index time series was found to be stationary. We conducted cointegration analysis to formally test for the existence of a long run relationship between information security firms' stock returns and malicious intensity. If the variables are cointegrated, a vector error correction model (VECM) is required. The VECM can be estimated by maximum likelihood using Johansen's cointegration approach [34] that is useful in detecting stationary linear combinations between variables. Johansen's cointegration analysis with intercepts, seasonal dummies, and no restrictions on the intercepts led to a cointegration order of zero. The zero cointegration order was confirmed by the Lambda max test [34] of the null hypothesis that there are zero cointegrated vectors against the alternative of one cointegrated vector. The zero cointegration order was also confirmed by the trace test [34], which pointed out that the null had at most zero cointegrated vectors, while the alternative had four cointegrated vectors (the determination of the number four was made using the Hannan Quinn information criterion [27]). These results suggest that a stationary linear combination does not exist between the two variables and that a VECM does not yield any additional information compared with a VAR model in first differences. We performed first order differencing on the malicious intensity time series. The differenced malicious intensity time series appeared stationary and showed no autocorre lation, which justified running further regression analyses to test the validity of our hypothesis. The VAR model of order p, following [52], is given in Eq. (1).

$$\mathrm{Y_t\ \ Y_{t-1}=C+A_1[Y_{t-1}\ \ Y_{t-2}]+A_2[Y_{t-2}\ \ Y_{t-3}]+\ \ldots+A_p\left[Y_{t-p}\ \ Y_{t-p-1}\right]+e_t} \tag{1}$$

where the vector $\mathrm{Y}_t = \begin{bmatrix} \text{Security return index}_t \\ \text{Malicious intensity}_t \\ \text{Market index}_t \end{bmatrix}$ is a 3×1 vector representing the daily stock market returns of information security firms, malicious intensity, and the market return at time t. C is a 3×1 vector of constants, $\mathrm{A_i}$'s are 3×3 matrices (for i = 1, …, p), and e_t is a 3×1 error vector (e'_t is its transpose) whose error terms have mean 0, a contemporaneous covariance matrix Ω, and no serial correlation i.e. $\mathrm{E(e_t)=0}$, $\mathrm{E(e_t,\ e'_t)=\Omega}$ and $\mathrm{E(e_t,\ e_{t-k})=0}$ for every non zero k. Given that the independent variables in the three equations, represented in Eq. (1), are all the same, there is no efficiency gain in using seemingly unrelated regression. Therefore, all three equations were estimated by the ordinary least squares (OLS) method.

4.1. VAR results

Table 3 gives the results of the VAR analysis upon relating the daily intensity of malicious attacks to the daily stock market returns of

Table 3
VAR results.

Explanatory variables	OLS estimate	t-value	Lag in days
VAR results for security return index			
$\mathrm{DIF_1}$[malicious intensity]	55.82758E 04	3.02**	30 (1 month)
Market index	50.12391E 02	2.18*	30 (1 month)
VAR results for $\mathrm{DIF_1}$[malicious intensity]			
Security return index	39.03055	2.37*	120 (4 months)
Market index	55.72448	2.92**	60 (2 months)
VAR results for market index			
Security return index	11.82465E 01	2.18*	60 (2 months)
	22.76156E 02	2.56*	90 (1 quarter)
$\mathrm{DIF_1}$[malicious intensity]	14.09541E 03	1.96*	90 (1 quarter)

**Significant at the 1% level.
*Significant at the 5% level.

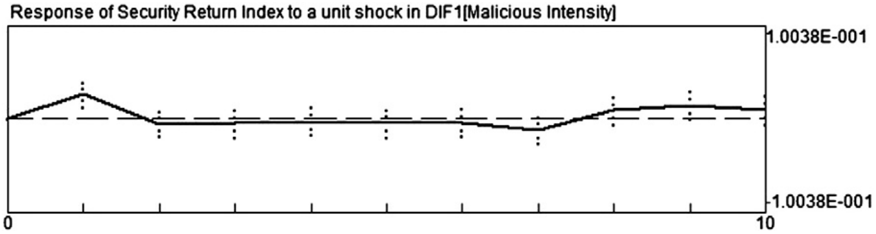

Fig. 5. Response of security return index to a unit shock in malicious intensity (~300-day horizon).

information security firms, while controlling for the market index. Joint significance tests confirmed that all included variables are indispensible. We found a significant Granger causality effect of malicious intensity on the stock market returns of information security firms (coefficient = 0.005583, p value < 0.01, Lag = 30 days). After controlling for the market index, a 1 month lagged increase in malicious intensity has a positive effect on the stock market returns of information security firms. In Table 3, we use DIF_i to denote ith order differencing. We further confirmed the significance of the selected lagged variables using the variance decomposition method, based on impulse response functions, which permits us to study the response of each variable to a unit shock in another variable. These impulse response functions study how a particular variable in the system responds to unit changes in another variable, while all other shocks are held at zero.

Fig. 5 represents the graph of the impulse response corresponding to our hypothesis, with 1% error bands. The x axes in the figures represent the forecast horizon for 10 months (approximately 300 days) forward. The y axes represent the response of a given dependent variable to a unit shock in a given independent variable, while holding constant the remaining variables. The two time standard error band corresponds to the point wise 99% confidence interval of each innovation response. When the two time standard error band is above (below) the horizontal axis, the response is significantly positive (negative). The results from Fig. 5 confirm, at the 1% significance level, a Granger causality effect of malicious intensity on the stock market returns of information security firms.

Other interesting results are also revealed in Table 3. Malicious intensity is shown to be negatively associated with prevailing market conditions (coefficient = −55.724480, p value < 0.01, Lag = 60 days). This is consistent with the results in [38] where it was shown that malicious attacks get more virulent in down markets. On the other hand, when the stock market is thriving, firms and individuals have more cash on hand to invest in information security, which helps in alleviating the negative effects of malicious attacks. Table 3 also shows that malicious attacks can have a long term negative effect on the stock market (coefficient = −0.014095, p value < 0.05, lag = 90 days). In fact, it has been shown that security breaches that could result from malicious attacks have a significant negative effect on the stock price of breached firms [8,23,30,40]. The overall negative effect of malicious attacks on the stock market indicates that their negative consequences on the stock price of breached firms far more outweighs their benefits vis à vis the market value of information security firms.

5. Time-delayed ANN-based analysis

Numerous studies have proposed ANNs as a viable and even better alternative than regression analysis for classification and prediction problems. ANNs have been widely used for financial forecasting [9,36,57,59,64 66]. The application of ANNs to information security has been used in the areas of intrusion detection and prevention, and spam detection and control [14]. To relate lagged versions of the variables, as was done for the VAR analysis in Section 4, we propose to use the time delayed ANN structure, a nonlinear time series model that is time lagged. Time delayed ANNs have been shown to be particularly useful for forecasting stock prices [68]. A time delayed ANN is based on a multilayer perceptron (MLP) neural network design in which the input is formed by a delayed segment of a time series [28]. The algorithm and MATLAB code for the time delayed ANN model implementation are given in Appendix C [31]. The inputs to our model are the daily aggregated intensity of malicious attacks and the daily levels of the NYSE/AMEX/NASDAQ value weighted index. The output is the security return index that is used as a proxy for the stock market performance of the information security market segment.

We divided the time series into a training sample to build the model and a testing sample to validate its performance. Although there are no clear cut guidelines on how to partition datasets into training and testing sub samples, Hair et al. [26] recommends a three to one split between the two samples. We correspondingly divided our data points into three quarters training data and one quarter testing data. The backpropagation (BP) algorithm is employed to train the network; this is because Cybenko [16] recognized the power of the BP algorithm to approximate any continuous function using a feed forward network of appropriate topology. To achieve backpropagation learning, the activation function needs to be both bounded and differentiable so that it can implement the underlying gradient search weight updating formula. We chose the logistic function as the activation function for the output layer, as recommenced in [28]. Although other considerations, such as choosing the epoch size (i.e. the size of the batch of data that are swept through during each training and testing run of the MLP neural network), learning rate, momentum term, termination criteria, among others, are necessary for the software implementation of the MLP neural network, this paper only discusses a select set of MLP design and implementation issues relevant to our problem. We refer the interested reader to [28] for an extensive and more general coverage of the subject. Many researchers have proposed different heuristic methods to determine the required number of neurons for optimal ANN performance. Haykin [28] conjectured that N input patterns require $(N-1)$ neurons in the single hidden layer network case and argued that simple networks have more effective generalizing capabilities than more complex networks do with more hidden layers. Smaller networks learn and operate more quickly [25] because the shortage of units forces the algorithm to develop general rules to discriminate between input patterns, whereas it would otherwise tend to learn each item of data as a special case [17]. Accordingly, we selected the simplest MLP topology to avoid over fitting. We used two metrics, namely adjusted R^2 and the mean square error (MSE), to evaluate a particular MLP topology.

5.1. Comparative results

The time delayed ANN topology with only one hidden layer and four neurons gave a relative predictive accuracy of 95%, 93.44%

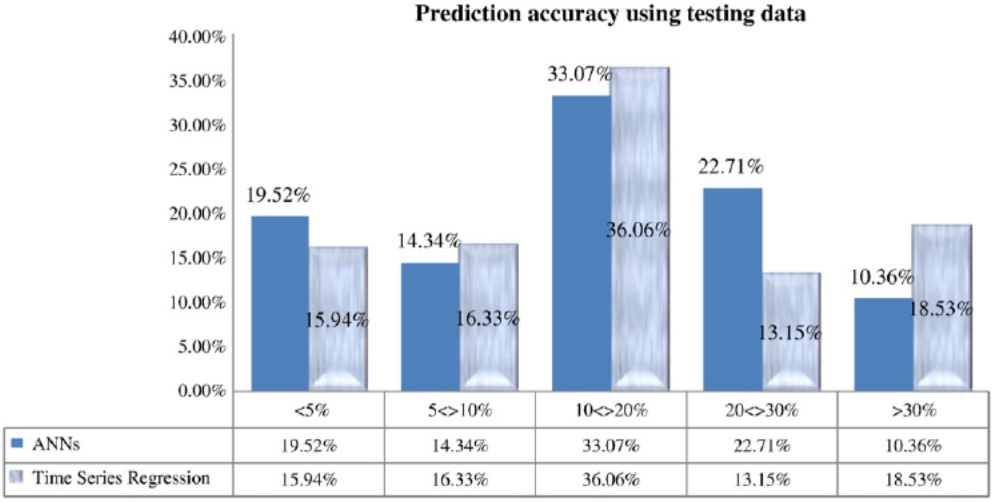

	<5%	5<>10%	10<>20%	20<>30%	>30%
■ ANNs	19.52%	14.34%	33.07%	22.71%	10.36%
▨ Time Series Regression	15.94%	16.33%	36.06%	13.15%	18.53%

Fig. 6. Predictive accuracy on testing data.

adjusted R^2, and a near zero mean square error. Comparatively, VAR analysis relating daily values gave an adjusted R^2 of 72.92% (22% lower than ANN) and a relative predictive accuracy of 84.98% (10.5% lower than ANNs). We tested whether the ANN's outputs differed significantly from those of VAR analysis. We conducted a paired two sample for means *t* test with one quarter of the data, i.e. 819 testing data points, and obtained a t value of 7.1319, which is much larger than the critical t value of 2.586 at the 99% significance level, confirming that both outputs are significantly different. Fig. 6 gives the percentages of the ANN's predicted responses that fall within 5% and 30% of the actual output. Approximately 20% of the time delayed ANN's predicted outputs fell within less than 5% (compared with less than 16% for VAR) and only about 11% were above 30% (compared to more than 18% for VAR) of the actual outputs.

5.2. ANN sensitivity analyses

1. Varying the ANN topology. We studied the effects of varying the number of neurons and hidden layers in the time delayed ANN. Fig. 7 shows the predictive accuracy of the ANN, with a varying number of neurons, given one hidden layer. As the number of neurons increases, the ANN's predictive accuracy increases accordingly until a certain threshold, after which the added complexity starts to have an adverse effect. Fig. 8 plots the ANN's predictive accuracy for a varying number of layers, given four neurons per layer. The predictive accuracy is relatively high up to five layers. With six layers and beyond, however, the predictive accuracy sharply falls to zero as a result of over fitting. This means that a network that is too complex may fit the noise, not just the signal. Over fitting is problematic because it can lead to predictions that are far beyond the range of the training data, especially in MLP based ANN implementations.
2. Varying data quality. Bansal et al. [4] argued that data quality can be described in terms of a number of dimensions, including frequency, accuracy, and response time. In general, a model performs better the more frequent and the more accurate the data. They showed in their paper that neural network based forecasts were more robust than linear regression forecasts when data accuracy decreased. We compared the performance of the designed MLP neural network with various data frequencies, i.e. daily and monthly intensity and price data. The main metric we used for comparison is predictive accuracy. Before proceeding with the ANN results, we compared the predictive accuracy of VAR analysis using monthly and daily data. In the case of monthly data, the predictive accuracy amounted to 82.47%, while in the case of daily data the predictive accuracy was 84.98%. We subsequently ran the previous ANN analysis with 32 different configurations ranging from one hidden layer and two neurons to four hidden layers and nine neurons. Fig. 9, a scatter plot showing the daily and monthly predictive accuracies, reveals that daily accuracies are

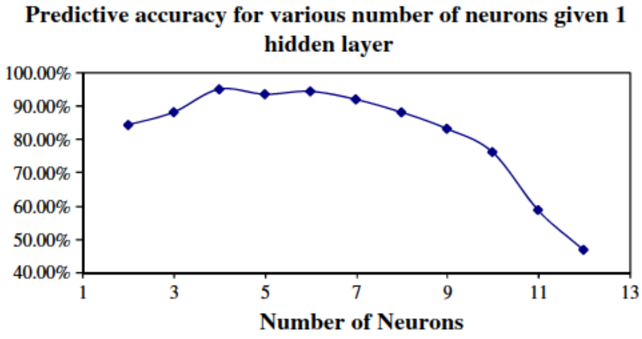

Fig. 7. Predictive accuracy given 1 hidden layer.

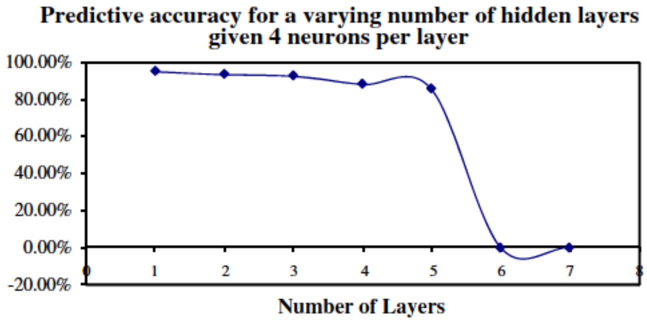

Fig. 8. Predictive accuracy at 4 neurons per layer.

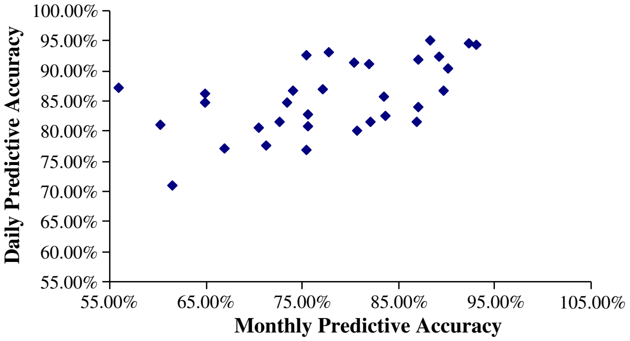

Fig. 9. Scatter plot of daily and monthly predictive accuracies.

consistently higher than their monthly counterparts. The results of the paired *t* test between the daily and monthly predictive accuracies confirmed that daily and monthly accuracies are significantly different and that the daily predictive accuracy is significantly higher (The obtained t value of 5.21 is much higher than the critical t value of 2.04). This is expected as more aggregation in the context of data mining leads to loss of important information.

6. Discussion and conclusion

6.1. Summary of results and implications

We had two key objectives in this research: (i) to study the effect of malicious attacks on the stock market returns of information security firms; and (ii) to examine the value in using complementary methods such as artificial neural networks (ANNs) and VAR models for predicting the stock price performance of information security firms in the presence of malicious attacks. Our analyses indicated that malicious intensity has a 1 month lagged positive effect on the stock market returns of information security firms. We also showed that malicious intensity follows an upward trend in down markets, a result that is consistent with the findings in [38]. Although the market returns of information security firms increase with the intensity of malicious attacks, the resulting security and privacy breaches and their damaging effects on breached firms are overpowering and have been shown [8,23,30,40] to drag the overall stock market down. Our results confirmed these findings. On the other hand, our results also revealed that an upward trending stock market, signaling that firms are profitable and are able to invest in information security, does eventually reduce the intensity of malicious attacks, as shown in [38].

The results imply that stock market investors pay close attention to the overall intensity of malicious attacks, thus highlighting the importance of information security to protect consumers of IT, whether firms or individuals. As long as IT products are vulnerable to malicious attacks, stock market investors expect IT consumers to invest in information security to get the needed protection against malicious attacks. The results also imply that investors blame malicious attacks on limited or ineffective investment in information security. In fact, if stock market investors did not expect malicious attacks to trigger an increase in information security investments, malicious attacks would not have positively affected the stock market returns of information security firms significantly. Our results also suggest that studying the intensity of malicious attacks is profitable to stock market investors interested in owning shares in the security market sector.

6.2. Limitations

This study, like any time series related study, has its limitations. First, as is typical with any aggregated data, some data loss is unavoidable. However, malicious attacks were only aggregated daily, rather than monthly or annually, so the loss of information is minimal. We could not think of a better approach since the smallest granularity of stock prices is daily. It is appropriate to note at this point that event study analyses that have widely studied the effect of security breaches on the stock price of firms [8,10,22,23,30,40] also aggregate the data over days as well as over firms to compute the cumulative abnormal returns and assess their significance. Secondly, the reader should be cautious about generalizing the results of this study beyond the 1996 2008 time period. Although this period is quite diverse in terms of the economic events it encompasses (e.g. pre bubble period, September 11), it cannot possibly be a perfect depiction of the future. The analyses and results should be updated as new malicious and stock market data become readily available.

6.3. Theoretical, methodological, and practical contributions

Information security breaches that result from malicious attacks, have been shown to positively impact the stock price of information security firms using event study analyses [10,22]. Event studies consist of comparing the expected returns for a firm's stock, assuming the event did not occur, with the actual returns after the event has occurred. Their time span is, however, short, merely extending over few days around the announcement of the event. Time series analyses, on the other hand, depict patterns and study relationships over time. Using time series analyses, thus, allowed us to capture the extent of the damages of malicious attacks on firms, and study the time lag effects between variables. Extracting malicious and stock market data and developing time series out of them constitute a novel attempt that sets methodological and theoretical precedence for future research.

Unlike risk management and RO based analyses [37] that impose assumptions about the distribution of the data, the analyses in this paper use real data pertaining to actual malicious attacks and publicly available stock market data. Although secondary in nature, stock market data are available from reliable financial databases that are tightly regulated by the SEC. The stock market data from CRSP have been heavily used in prior research and deemed reliable. As far as malicious intensity data are concerned, breached firms are unwilling to share information about the extent of the damages inflicted upon their IT systems. They are either unable to compute an accurate dollar value of these damages or they are reluctant to share sensitive information that exposes the weaknesses in their systems. For either reason, data pertaining to malicious intensity are not easy to obtain. We made use of malicious attacks data that are collected and assessed by subject matter experts working for Symantec, the leader in antivirus technology. The way we computed the intensity of malicious attacks by combining the severity and number of malicious attacks is also novel and adept. Since we are only interested in the effect of malicious attacks on the stock market returns of information security firms, what really matters are the revenues that these firms accumulate as a result of malicious attacks. We weighed the number of malicious attacks by their severity levels, which would make more sense than accounting for either separately from the viewpoint of information security firms.

We have shown that our time delayed ANN model serves as a complementary approach to conventional VAR analysis. VAR analysis helped us establish the significance of the Granger causality relationship between the intensity of malicious attacks and the stock market returns of information security firms and determine the time lag and, thus, specify the inputs to the ANN model. Our time delayed ANN model resulted in 95% predictive accuracy, compared to around 85% for the regression counterpart.

The time delayed ANN implementation can be further compiled into an Excel plug in using the MATLAB Excel builder (http://www.mathworks.com). The resulting decision support system serves as an investment tool for mutual funds, hedge funds, and other investors who wish to balance their portfolios and protect them against stock price fluctuations resulting from security breaches. For example, hedge funds investing in information retrieval services companies, such as Amazon, Yahoo!, Ebay, and Google are at risk of losing significant market value upon the occurrence of malicious attacks and their resulting security breaches, which are especially debilitating for these information service providers. These hedge funds can use the proposed time delayed ANN model and decision support system to determine how much to invest in the information security market sector, thereby reducing the risk of losing money as a result of malicious attacks.

6.4. Avenues for future work

As future work and to gain even more accurate and statistically sound predictions, we propose to augment the time delayed ANNs with genetic algorithms (GAs), whose unstructured search and decision trees assist ANNs in the task of variable selection. GAs can be used to better select the time delayed ANN architecture including the number of hidden layers and number of hidden neurons in every layer. GAs can also be used to determine the variables from the input layer that have the highest correlation with the stock return output. As such more characteristics of malicious attacks, including their type (e.g. Trojans, worms) and origin, can be accounted for in the model. We also plan to incorporate attackers, firms, and other meaningful actors into the framework to add more realism and flexibility to our analyses.

We believe this paper raised interesting research issues through the proposed measures, methodologies, model, and findings that will help other researchers make sense of quantifiable malicious intensity and widely available financial data.

Acknowledgements

The authors would like to thank Professor Yu Hen Hu, whose guidance on an earlier project provided motivation for this research. We would also like to thank the Editors, especially Professor James Marsden, and the three anonymous reviewers for their valuable comments and suggestions.

Appendix A

Table A.1
List of public information security firms (Stock ticker in parentheses).

Public antivirus (AV) firms	Public network security firms		Public identity and access management (IAM) firms	
CyberGuard Corp (CGFW)	Applied Theory Corp (ATHY)	Trusted Information Systems (TISX)	ActivCard Corp (ACTI)	L-1 Identity Solutions Inc (ID)
eLinear Inc (ELU)	Cavium Networks Inc (CAVM)	Verint Systems Inc (VRNT)	Adaptive Solutions Inc (ADSO)	Liska Biometry Inc (LSKA)
Guardian Technologies (GRDN)	C-COR Electronics Inc (CCBL)	WatchGuard Technologies Inc (WGRD)	Alien Technology Corp (RFID)	Macrovision Corp (MVSN)
McAfee Inc (MFE)	Checkpoint Software (CHKP)	Websense (WBSN)	Applied Digital Solutions Inc (ADSX)	Netegrity Inc (NETE)
Procera Networks Inc (PKT)	Citrix Systems (CTXS)		AuthenTec Inc (AUTH)	NetIQ Corp (NTIQ)
Symantec Corp (SYMC)	Cylink Corp (CYLK)		AXENT Technologies Inc (AXNT)	Novell Inc (NOVL)
	Diversified Security Solutions (DVS)		Bindview Development Corp (BVEW)	Phoenix Technologies Ltd (PTEC)
	E-Biz Solutions Inc (EBIZ)		Blue Coat Systems Inc (BCSI)	RSA Security Inc (RSAS)
	EFJ Inc (EFJI)		Cogent Inc (COGT)	Saflink Corp (SFLK)
	eSoft Inc (ESFT)		Cognizant Tech Solutions Corp (CTSH)	Security Dynamics Technologies (SDTI)
	F5 Networks (FFIV)		Commun Intelligence Corp (CICI)	Tegal Corp (TGAL)
	GRIC Communications Inc (GRIC)		Convera Corp (CNVR)	Transcrypt International Inc (TRII)
	Info Resource Engineering Inc (IREG)		CyberSource Corp (CYBS)	TSL Inc (TSLI)
	International Network Services (INSS)		Digital Angel Corp (DOC)	TTR Technologies Inc (TTRE)
	Internet Sec Sys Inc (ISS)		Digital Biometrics Inc (DBII)	Tumbleweed Communications Corp (TMWD)
	ITI Technologies Inc (ITII)		Document Security Systems Inc (DMC)	VASCO Corp (VAS)
	Kanbay International (KBAY)		Double-Take Software Inc (DBTK)	VASCO Data Security Intl Inc (VDSI)
	Litronic Inc (LTNX)		Entrust Inc (ENTU)	VeriSign Inc (VRSN)
	NetLogic Microsystems (NETL)		ePresence Inc (EPRE)	Viisage Technology Inc (VISG)
	NetScreen Technologies Inc (NSCN)		First Data Corp (FDC)	WidePoint Corp (WYY)
	Network Associates Inc (NET)		I D Systems Inc (IDSY)	Zebra Technologies Corp (ZBRA)
	Network Engines (NENG)		Identix Inc (IDNX)	Zix Corp (ZIXI)
	NPS International Corp (NPSZ)		Imageware Systems Inc (IW)	
	Rainbow Technologies Inc (RNBO)		Intelli-Check Inc (IDN)	
	SafeNet Inc (SFNT)		InterTrust Technologies Corp (ITRU)	
	Secure Computing (SCUR)		iPass Inc (IPAS)	
	SonicWALL (SNWL)		IQ Biometrix Inc (IQBM)	
	Sourcefire Inc (FIRE)			
	SteelCloud inc (SCLD)			

Appendix B

Table B.1
Time series of number and intensity of malicious attacks.

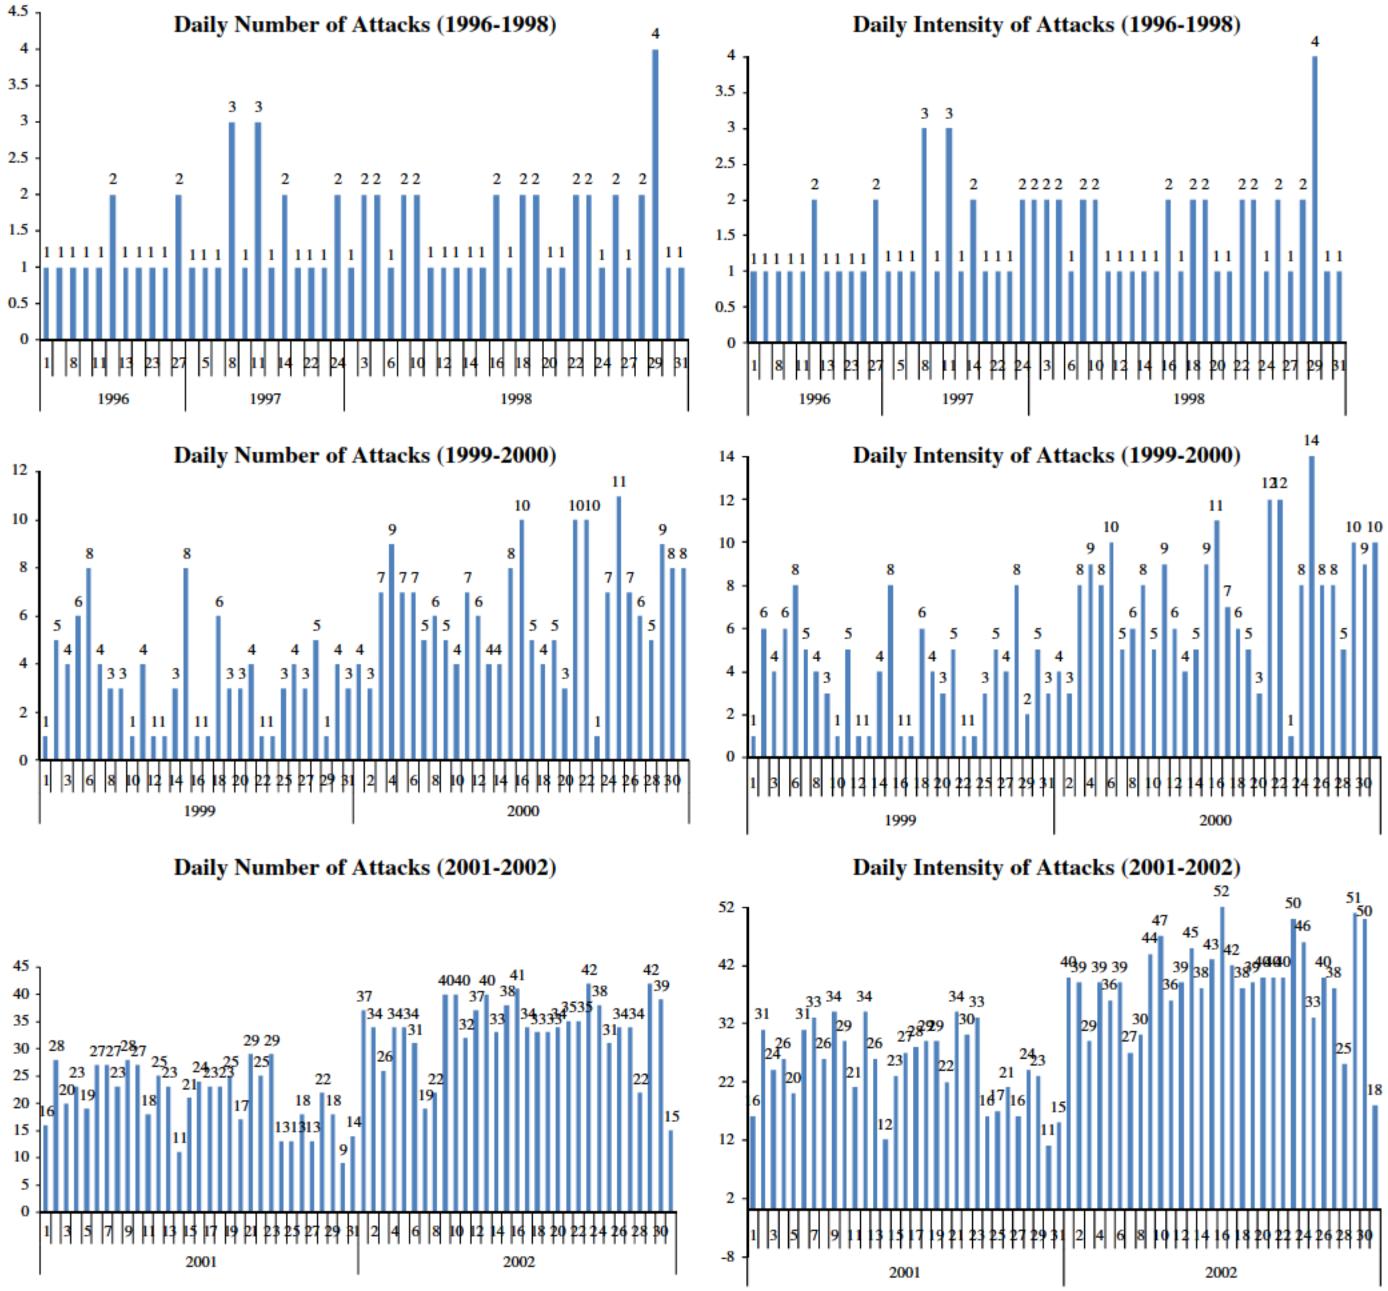

(continued on next page)

Table B.1 (*continued*)

Daily Number of Attacks (2003-2004)

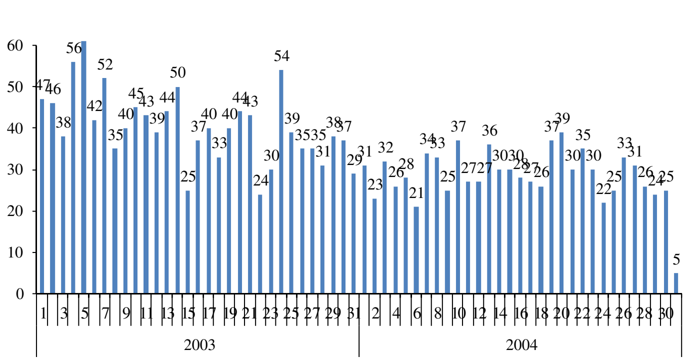

Daily Intensity of Attacks (2003-2004)

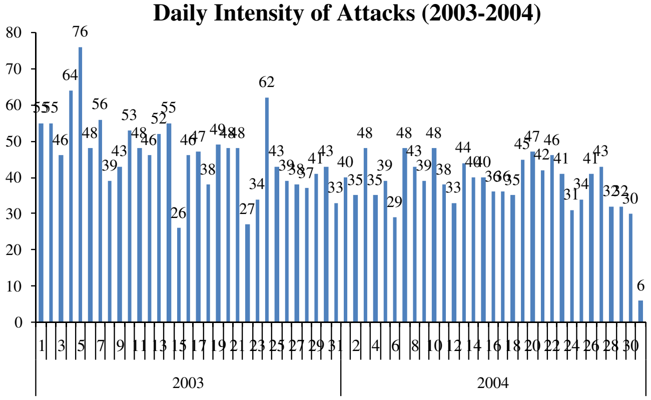

Daily Number of Attacks (2005-2006)

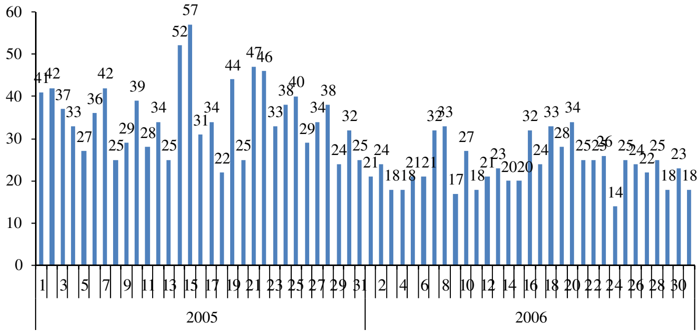

Daily Intensity of Attacks (2005-2006)

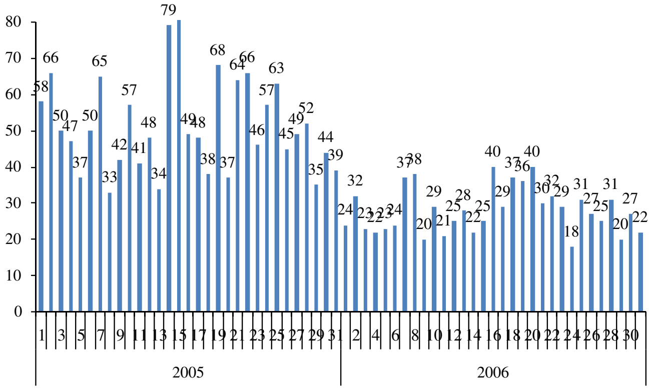

Daily Number of Attacks (2007-2008)

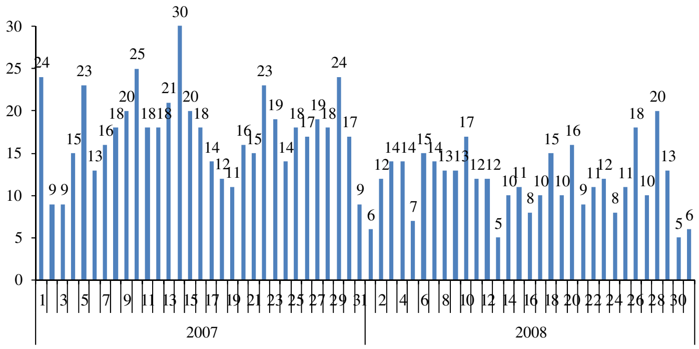

Daily Intensity of Attacks (2007-2008)

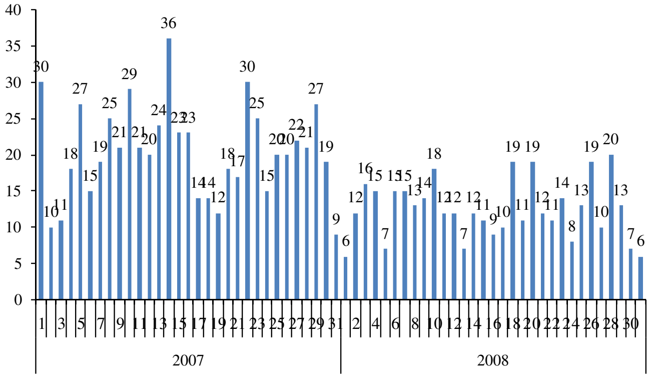

Appendix C

Time delayed ANN implementation

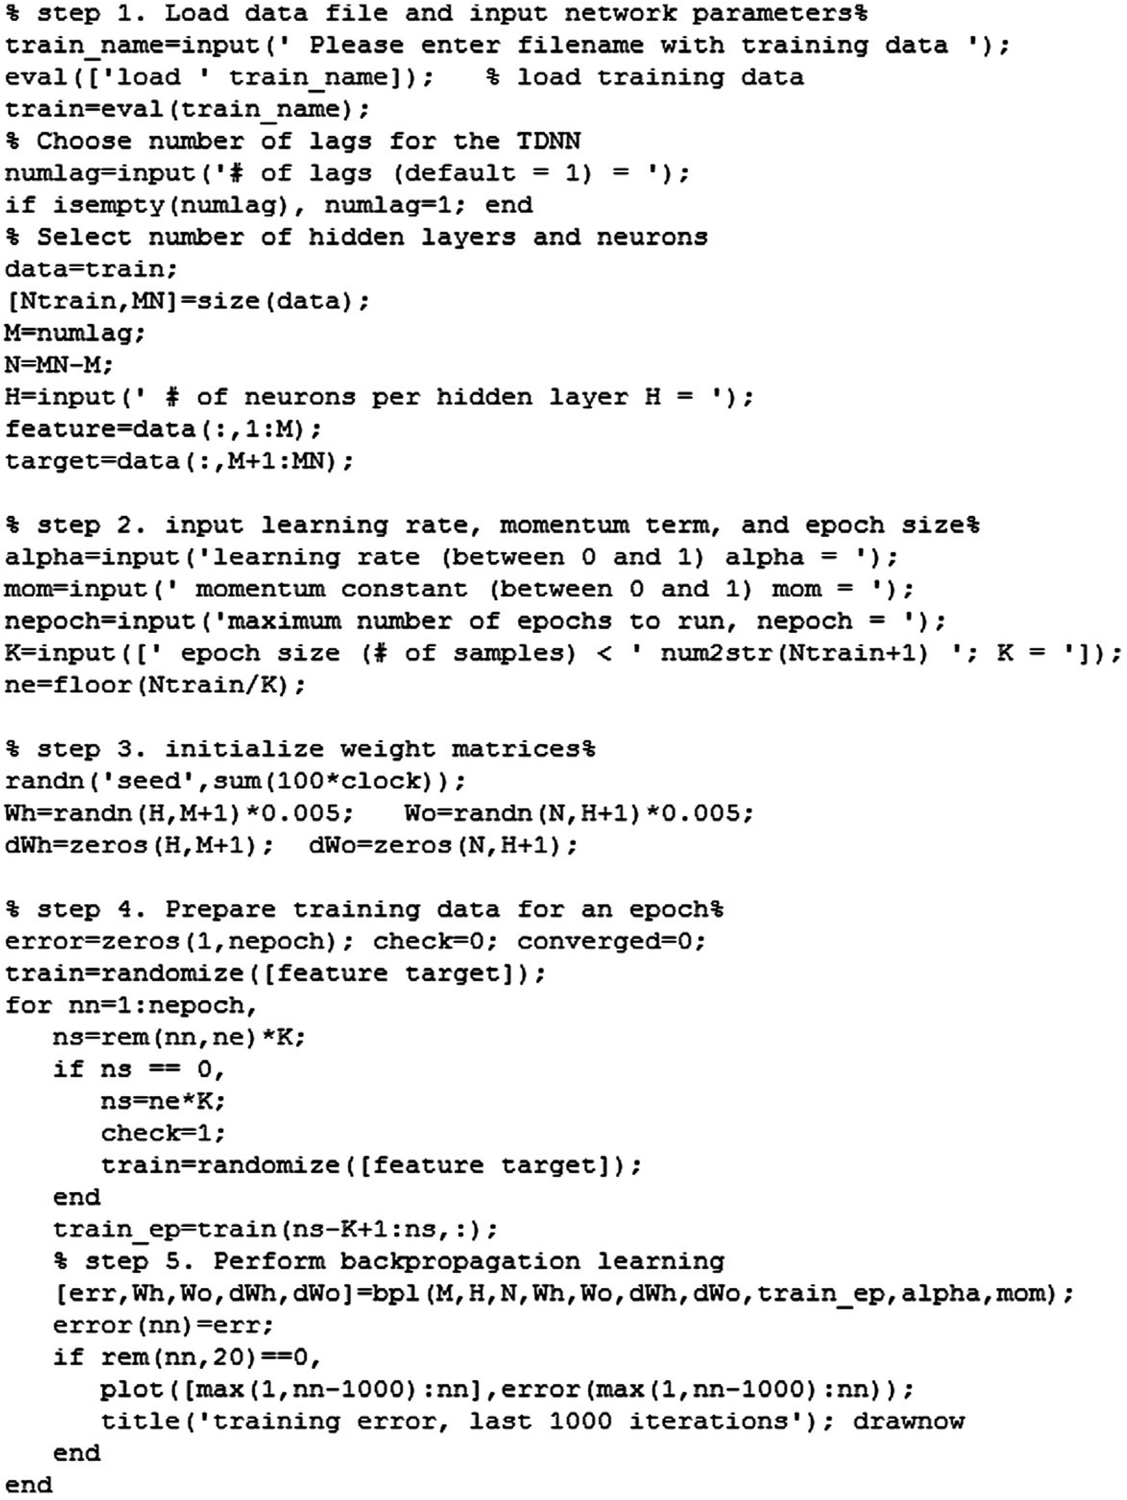

```
% step 1. Load data file and input network parameters%
train_name=input(' Please enter filename with training data ');
eval(['load ' train_name]);   % load training data
train=eval(train_name);
% Choose number of lags for the TDNN
numlag=input('# of lags (default = 1) = ');
if isempty(numlag), numlag=1; end
% Select number of hidden layers and neurons
data=train;
[Ntrain,MN]=size(data);
M=numlag;
N=MN-M;
H=input(' # of neurons per hidden layer H = ');
feature=data(:,1:M);
target=data(:,M+1:MN);

% step 2. input learning rate, momentum term, and epoch size%
alpha=input('learning rate (between 0 and 1) alpha = ');
mom=input(' momentum constant (between 0 and 1) mom = ');
nepoch=input('maximum number of epochs to run, nepoch = ');
K=input([' epoch size (# of samples) < ' num2str(Ntrain+1) '; K = ']);
ne=floor(Ntrain/K);

% step 3. initialize weight matrices%
randn('seed',sum(100*clock));
Wh=randn(H,M+1)*0.005;   Wo=randn(N,H+1)*0.005;
dWh=zeros(H,M+1);  dWo=zeros(N,H+1);

% step 4. Prepare training data for an epoch%
error=zeros(1,nepoch); check=0; converged=0;
train=randomize([feature target]);
for nn=1:nepoch,
   ns=rem(nn,ne)*K;
   if ns == 0,
      ns=ne*K;
      check=1;
      train=randomize([feature target]);
   end
   train_ep=train(ns-K+1:ns,:);
   % step 5. Perform backpropagation learning
   [err,Wh,Wo,dWh,dWo]=bpl(M,H,N,Wh,Wo,dWh,dWo,train_ep,alpha,mom);
   error(nn)=err;
   if rem(nn,20)==0,
      plot([max(1,nn-1000):nn],error(max(1,nn-1000):nn));
      title('training error, last 1000 iterations'); drawnow
   end
end
```

References

[1] E. Alfaro, N. Garcia, M. Gamez, D. Elizondo, Bankruptcy forecasting: an empirical comparison of AdaBoost and neural networks, Decision Support Systems 45 (1) (2008) 110–122.

[2] R. Anderson, Why information security is hard — an economic perspective, Proceedings of 17th Annual Computer Security Applications Conference, New Orleans, Louisiana, 2001.

[3] A. Arora, R. Krishnan, R. Telang, Y. Yang, An empirical analysis of software vendors' patch release behavior: impact of vulnerability disclosure, Information Systems Research 21 (2010) 115–132.

[4] A. Bansal, R.J. Kauffman, R.R. Weitz, Comparing the modeling performance of regression and neural networks as data quality varies: a business value approach, Journal of Management Information Systems 10 (1993) 11–32.

[5] B.M. Barber, M.N. Darrough, Product reliability and firm value: the experience of American and Japanese automakers, 1973–1992, Journal of Political Economy 104 (5) (1996) 1084.

[6] A. Blumstein, J. Cohen, D. Nagin, Introduction in Deterrence and Incapacitation: Estimating the Effects of Criminal Sanctions on Crime Rates, National Academy of Sciences, Washington, D.C., 1978

[7] M.A. Boyacioğlu, Y. Kara, O.K. Baykan, Predicting bank financial failures using neural networks, support vector machines and multivariate statistical methods: a comparative analysis in the sample of savings deposit insurance fund (SDIF) transferred banks in Turkey, Expert Systems with Applications 36 (2) (2009) 3355–3366.

[8] K. Campbell, L.A. Gordon, M.P. Loeb, L. Zhou, The economic cost of publicly announced information security breaches: empirical evidence from the stock market, Journal of Computer Security 11 (2003) 431–448.

[9] Q. Cao, M.E. Parry, Neural network earnings per share forecasting models: a comparison of backward propagation and the genetic algorithm, Decision Support Systems 47 (1) (2009) 32–41.

[10] H. Cavusoglu, B. Mishra, S. Raghunathan, The effect of Internet security breach announcements on market value: capital market reactions for breached firms and Internet security developers, International Journal of Electronic Commerce 9 (2004) 69–104.

[11] H. Cavusoglu, S. Raghunathan, W. Yue, Decision-theoretic and game-theoretic approaches to it security investment, Journal of Management Information Systems 25 (2) (2008) 281–304.

[12] Y. Chen, S. Ganesan, Y. Liu, Does a firm's product recall strategy affect its financial value? An examination of strategic alternatives during product-harm crises, Journal of Marketing 73 (6) (2009) 214–226.

[13] W.K. Chiang, D. Zhang, L. Zhou, Predicting and explaining patronage behavior toward web and traditional stores using neural networks: a comparative analysis with logistic regression, Decision Support Systems 41 (2) (2006) 514–531.

[14] C. Chou, A.P. Sinha, H. Zhao, Commercial Internet filters: perils and opportunities, Decision Support Systems 48 (4) (2010) 521–530.

[15] E.K. Clemons, Evaluation of strategic investments in information technology, Communications of the ACM 34 (1) (1991) 22–36.

[16] G. Cybenko, Approximation by superpositions of a sigmoidal function, Math, Control, Signals, and Systems 2 (1989) 303–314.

[17] K.B. DeTienne, D.H. DeTienne, S.A. Joshi, Neural networks as statistical tools for business researchers, Organizational Research Methods 6 (2003) 236–265.

[18] L. Ehrlich, Participation in illegitimate activities: a theoretical and empirical investigation, Journal of Political Economy 81 (1973) 521–564.

[19] E.F. Fama, Efficient capital markets: a review of theory and empirical work, Journal of Finance 25 (2) (1970) 383–417.

[20] E.F. Fama, Efficient capital-markets 2, Journal of Finance 46 (5) (1991) 1575–1617.

[21] P. Fraser, N. Groenewold, US share prices and real supply and demand shocks, The Quarterly Review of Economics and Finance 46 (1) (2006) 149–167.

[22] A. Garg, J. Curtis, H. Halper, Quantifying the financial impact of IT security breaches, Information Management & Computer Security 11 (2003) 74–83.

[23] S. Goel, H. Shawky, Estimating the market impact of security breach announcements on firm values, Information and Management 46 (7) (2009) 404–410.

[24] C.W.J. Granger, P. Newbold, Forecasting Economic Time Series, 2nd ed, Academic Press, 1986.

[25] M. Hagiwara, Removal of hidden units and weights for backpropagation networks, Proceedings 1993 Int. Joint Conf. Neural Networks, 1, 1993, pp. 351–354.

[26] J.F. Hair, R.E. Anderson, R.L. Tatham, W.C. Black, Multivariate Data Analysis, Prentice Hall, New Jersey, 1998.

[27] E.J. Hannan, B.G. Quinn, The determination of the order of an autoregression, Journal of the Royal Statistical Society 41 (1979) 190–195.

[28] S. Haykin, Neural networks: A Comprehensive Foundation, Macmillan Publishing, New York, 1994.

[29] K.B. Hendricks, V.R. Singhal, The effect of supply chain glitches on shareholder wealth, Journal of Operations Management 21 (2003) 501–522.

[30] A.J. Hovav, J. D'Arcy, The impact of denial-of-service attack announcements of the market value of firms, Risk Management and Insurance Review 6 (2) (2003) 97–121.

[31] Y.H. Hu, Time series analysis, http://homepages.cae.wisc.edu/~ece539/matlab/index.html, (Last accessed April 29, 2010).

[32] C. Isidore, Toyota recall costs: $2 billion, http://money.cnn.com/2010/02/04/news/companies/toyota_earnings.cnnw/?postversion=2010020410, (Last accessed April 29, 2010).

[33] G. Jarrell, S. Peltzman, The impact of product recalls on the wealth of sellers, The Journal of Political Economy 93 (1) (1985) 512–536.

[34] S. Johansen, Statistical analysis of cointegrating vectors, Journal of Economic Dynamics and Control 12 (1988) 231–254.

[35] J. Kamp, T.M. Burton, Corporate news: recalls hurt Medtronic earnings, Wall Street Journal, B.4, http://online.wsj.com/article/SB124272809096834003.html 2009, (Last accessed April 29, 2010).

[36] A. Keles, M. Kolcak, A. Keles, Adaptive neuro-fuzzy model for forecastingthe domestic debt, Knowledge-Based Systems 21 (2008) 951–957.

[37] L. Khansa, D. Liginlal, Valuing the flexibility of investing in security process innovations, European Journal of Operational Research 192 (2009) 216–235.

[38] L. Khansa, D. Liginlal, Quantifying the benefits of investing in information security, Communications of the ACM 52 (2009) 113–117.

[39] W.R. King, G. Premkumar, K. Ramamurthy, An evaluation of the role and performance of a decision support system in business education, Decision Sciences 21 (3) (2007) 642–659.

[40] D. Liginlal, I. Sim, L. Khansa, How significant is human error as a cause to privacy breaches? An empirical study and a framework for error management, Computers and Security 28 (3–4) (2009) 215–228.

[41] C.-J. Lu, T.-S. Lee, C.-C. Chiu, Financial time series forecasting using independent component analysis and support vector regression, Decision Support Systems 47 (2009) 115–125.

[42] Z. Meng, S.T. Lee, The value of IT to firms in a developing country in the catch-up process: an empirical comparison of China and the United States, Decision Support Systems 43 (3) (2007) 737–745.

[43] J. Muntermann, Towards ubiquitous information supply for individual investors: a decision support system design, Decision Support Systems 47 (2) (2009) 82–92.

[44] K. Palepu, Diversification strategy, profit performance and the entropy measure, Strategic Management Journal 6 (1985) 239–255.

[45] M. Paliwal, U.A. Kumar, Neural networks and statistical techniques: a review of applications, Expert Systems with Applications 36 (2009) 2–17.

[46] I. Pankratz, Forecasting with Dynamic Regression Models, Wiley, New York, 1991.

[47] G. Rasch, On general laws and the meaning of measurement in psychology, Proceedings of the 4th Berkeley Symposium on Mathematical Statistics and Probability, IV, University of Chicago Press, Berkeley, 1980, pp. 321–334.

[48] M. Rhee, P.R. Haunschild, The liability of good reputation: a study of product recalls in the U.S. automobile industry, Organization Science 17 (1) (2006) 101–117.

[49] A. Rombel, Internet security in an insecure world, Global Finance 15 (13) (2001) 28–32.

[50] P.H. Rubin, R.D. Murphy, G. Jarrell, Risky products, risky stocks, Regulation 12 (1) (1988) 35–39.

[51] N.G. Rupp, Are government initiated recalls more damaging for shareholders? Evidence from automotive recalls 1973–1998, Economics Letters 71 (2001) 265–270.

[52] C.A. Sims, Macroeconomics and reality, Econometrica 48 (1980) 1–48.

[53] A.P. Sinha, J.H. May, Evaluating and tuning predictive data mining methods using receiver operating characteristic curves, Journal of Management Information Systems 21 (2004–5) 249–280.

[54] K.J. Soo Hoo, How much is enough? A risk-management approach to computer security, Ph.D. dissertation, Graduate School of Engineering, Stanford University, Stanford, CA, 2000

[55] K.E. Spier, Product safety, buybacks, and the post-sale duty to warn, Journal of Law, Economics, and Organization, in press.

[56] W. Stallings, Network Security Essentials: Applications and Standards, 3rd ed. Prentice Hall, New Jersey, 2007.

[57] J. Sun, H. Li, Data mining for listed companies' financial distress prediction, Knowledge-Based Systems 21 (1) (2008) 1–5.

[58] L.L. Thurstone, A law of comparative judgment, Psychological Review 34 (1927) 278–286.

[59] E. Tsang, P. Yung, J. Li, EDDIE-Automation, a decision support tool for financial forecasting, Decision Support Systems 37 (4) (2004) 559–565.

[60] A. Vellidoa, P.J.G. Lisboaa, J. Vaughanb, Neural networks in business: a survey of applications (1992–1998), Expert Systems with Applications 17 (1999) 51–70.

[61] S. Walczak, T. Sincich, A comparative analysis of regression and neural networks for university admissions, Information Sciences 119 (1999) 1–20.

[62] H. Wang, A.S. Weigend, Data mining for financial decision making, Decision Support Systems 37 (4) (2004) 457–460.

[63] T. White, R. Pomponi, Gain a competitive edge by preventing recalls, Quality Progress 36 (8) (2003) 41–49.

[64] L. Yu, S. Wang, K.K. Lai, Credit risk assessment with a multistage neural network ensemble learning approach, Expert Systems with Applications 34 (2) (2008) 1434–1444.

[65] L. Yu, S. Wang, K.K. Lai, An intelligent-agent-based fuzzy group decision making model for financial multicriteria decision support: the case of credit scoring, European Journal of Operational Research 195 (3) (2009) 942–959.

[66] W. Zhang, Q. Cao, M.J. Schniederjans, Neural network earnings per share forecasting models: a comparative analysis of alternative methods, Decision Sciences 35 (2) (2004) 205–237.

[67] L. Zhou, J.K. Burgoon, D.P. Twitchell, T. Qin, J.A. Nunamaker, Comparison of classification methods for predicting deception in computer-mediated communication, Journal of Management Information Systems 20 (2004) 139–165.

[68] X. Zhu, H. Wang, L. Xu, H. Li, Predicting stock index increments by neural networks: the role of trading volume under different horizons, Expert Systems with Applications 34 (4) (2008) 3043–3054.

Lara Khansa is an Assistant Professor in the Department of Business Information Technology at Virginia Tech. She received her Ph.D. in Information Systems, M.S in Computer Engineering, and MBA in Finance and Investment Banking from the University of Wisconsin, Madison. Her primary research interests include the economics of information systems and information security, economic-based decision making in information systems, theories of investment, innovation, and market value creation in the particular realm of information systems, business value of information systems, and regulatory economics and its implications for IT innovation and the reshaping of the IT industry landscape. Dr. Khansa has published in leading journals including Decision Support Systems, European Journal of Operational Research, Communications of the ACM, Computers and Security, IEEE Technology and Society, and many others. She is a member of the Association for Information Systems (AIS), the Institute of Electrical and Electronics Engineers (IEEE), and the Beta Gamma Sigma National Honor Society.

Divakaran Liginlal is an Associate Teaching Professor of Information Systems at Carnegie Mellon University. Lal holds a B.S. in Telecommunication Engineering from CET, a M.S. in Computer Science from the Indian Institute of Science, and a Ph.D. in Management Information Systems from the University of Arizona. Before joining CMU, he taught at three U.S. universities, including nine years at the University of Wisconsin-Madison. Lal's research in information security and decision support systems has been published in such journals as CACM, IEEE-TKDE, IEEE-SMC, European Journal of Operational Research, Computers & Security, Decision Support Systems, and Fuzzy Sets & Systems. Lal's teaching and research have been supported by organizations such as Microsoft Corporation, Hewlett Packard, CISCO, and the ICAIR at the University of Florida.