

## Research Note

# A Value-at-Risk Approach to Information Security Investment

Jingguo Wang

College of Business Administration, University of Texas at Arlington, Arlington, Texas 76019, jwang@uta.edu

Aby Chaudhury

Bryant University, Smithfield, Rhode Island 02917, achaudhu@bryant.edu

H. Raghav Rao

School of Management, State University of New York at Buffalo, Buffalo, New York 14260, mgmtrao@buffalo.edu

Information security investment has been getting increasing attention in recent years. Various methods have been proposed to determine the effective level of security investment. However, traditional expected value methods (such as annual loss expectancy) cannot fully characterize the information security risk confronted by organizations, considering some extremal yet perhaps relatively rare cases in which a security failure may be critical and cause high losses. In this research note we introduce the concept of value-at-risk to measure the risk of daily losses an organization faces due to security exploits and use extreme value analysis to quantitatively estimate the value at risk. We collect a set of internal daily activity data from a large financial institution in the northeast United States and then simulate its daily losses with information based on data snapshots and interviews with security managers at the institution. We illustrate our methods using these simulated daily losses. With this approach, decision makers can make a proper investment choice based on their own risk preference instead of pursuing a solution that minimizes only the expected cost.

*Key words:* information assurance; security investment; value-at-risk (VaR); extreme value analysis

*History:* Seungjin Whang, Senior Editor; Ram Gopal, Associate Editor. This paper was received on February 1, 2006, and was with the authors 6 months for 3 revisions.

## 1. Introduction

The importance of effective information security management has increased in recent years due to the increasing frequency and cost of security incidents (Gordon et al. 2005). While high-risk organizations may adopt security at any price, most commercial organizations have to consider the cost-benefit trade-off for such investment. In the Ernst & Young Global Information Security Survey (Ernst & Young 2003, 2004), budget constraints are listed as one of the main obstacles to effective information security. Quantification tools, if applied prudently, can assist in the anticipation and control of direct and indirect computer security costs (Mercuri 2003, Geer et al. 2003).

In this research we propose an approach incorporating the concept of value-at-risk (VaR) (Crouhy et al. 2001, Duffie and Pan 1997, Jorion 1997) for information security investment and further intro-

duce extreme value analysis (Gumbel 1958) to estimate the VaR of daily losses. VaR was originally developed as a tool for assessing the risk of financial assets, and it is widely used in financial engineering and insurance. It is a statistical measure of the risk associated with an investment or set of investments. Extreme value theory quantifies the stochastic behavior of a process at unusually large (or small) levels; it is concerned with probabilistic and statistical questions related to those extremely rare events. To our knowledge, this is the first application of VaR and extreme value theory to information security investment decisions.

We consider the scenario of threat in a firm (Anderson and Brackney 2004, Chinchani et al. 2005, Shaw et al. 1998). On a daily basis, a large firm experiences some activities (which may be referred as exploits) that result in incidents of denial of services

(DoS) to its employees. According to the senior manager of information security at a financial institute in the northeast United States, some typical exploits that result in denial of services involve deleting or changing objects such as user objects, workstation objects, network components, and server components and subcomponents. Most of these incidents affect individual productivity in various ways, but a few of them may interrupt network services. Such disruptions affect normal business activities and cause a loss of productivity and reputation. In general, these activities occur on internal machines. They may be initiated by internal staff either accidentally or purposely, or manipulated by outsiders through malicious tools such as Trojan horses and other programs.

We regard an activity as having a particular probability of causing an incident. The occurrence probability and the cost of an incident are affected by security investments. A firm makes its investment choice based on its risk preference. On one hand, the firm may wish to minimize its expected daily cost, which includes the daily loss caused by incidents and the daily cost of countermeasures. On the other hand, the firm may expect that the daily loss caused by security incidents should not exceed a certain value, say \$10,000, with a certain level of confidence (for example, 99%). We use extreme value analysis to characterize the extreme (large) behavior of daily losses and estimate their VaR. With the application of the extreme value theory, we attempt to address the following issues:

1. What is the probability distribution of the high daily losses—i.e., what is the occurrence probability of daily loss above a given level? In other words, given a period (say, 100 days), what is the level that the daily loss is expected to exceed once in that period?

2. Are there seasonal phenomena or time trends in the extreme behavior of daily losses?

By answering these questions, we make inferences about the VaR and determine effective security investment levels to modulate the risk.

The contributions of this paper are threefold. First, we introduce the concept of VaR in the context of risk characterization for information security. Second, we use extreme value analysis to characterize daily losses and estimate their VaR. Finally, we demonstrate the approach in a scenario of protecting information systems against threats arising internally at a financial

institute. Decision makers can use this approach to make proper investment choices based on their own risk preferences instead of pursuing a solution that minimizes only the expected cost.

The organization of this paper is as follows: In §2, after a literature review, we introduce the concept of VaR for information security investment. In §3 we describe the daily internal activity data that we collected from a financial institute. In §4 we perform extreme value analyses on the simulated daily losses, estimating the extreme value model with and without the assumption that the observations are temporally independent. Seasonal phenomena and time trends in daily losses are examined, and a simple investment example is used to demonstrate how we can utilize the method. Finally, in §5 we summarize our study and discuss future research.

## 2. Security Risk and Security Investment

### 2.1. Related Literature

IT payoff has been widely examined (Devaraj and Kohli 2002), but research in IT security investment has gained interest only recently. The models that have been proposed to determine the effective level of security investment fall into basically two approaches: game theory to model strategic interactions between organizations and attackers, or traditional risk/decision analysis frameworks.

Some researchers have argued that information security can be treated as a game between organizations and attackers. While organizations try to cover vulnerabilities in their systems, attackers race to exploit them. Security investments not only prevent security incidents by reducing vulnerabilities that attackers can exploit, but also act as a deterrent for attackers by making attacks less attractive (Schechter and Smith 2003). Longstaff et al. (2000) argued that investment in system risk assessment can reduce the likelihood of intrusions, which yields benefits much higher than the investment. Varian (2004) examined the free rider problem in information security investment in different circumstances using a game theory model. In their models, both agents, the attacker and the defender, maximized their own expected benefits, and it was shown that the reactions of the defender

and attacker depend on their own cost–benefit ratio. Gal-or and Ghose (2005) studied economic incentives for sharing security information and found that security investments and security information sharing act as strategic complements in equilibrium. Kannan and Telang (2005) considered whether movement toward a market-based mechanism for vulnerability disclosure leads to a better social outcome.

However, the game theory approach suffers from the fact that the rationality of hackers is hard to capture in a model, because they may be motivated by a different value system. They may be rational, but it may not be on our terms; they may be driven by motivations other than money. It is hard for us to know their cost function for attacking the system.

Generally, traditional risk or decision analysis models apply a standard result in optimal-control theoretic certainty equivalence, which implies that only the mean values (probability-weighted average outcomes) of target variables matter for an optimal policy setting. Gordon and Loeb (2002) proposed one such model for the expected benefits of investment in information security. Hoo (2000) used a decision analysis approach to evaluate different policies for information security. Longstaff et al. (2000) proposed a hierarchical holographic model (HHM) to assess security risks and provide a model for assessing the efficacy of risk management.

Our model belongs to the decision theory approach. A risk situation in a decision theoretic model is usually characterized in terms of the expected value of a stochastic variable, such as annual loss expectancy (ALE), because in most applications the primary interest is in the center of the distribution, not the tail. Such methods cannot fully characterize the information security risk faced by a firm, however, because there are some extremal but relatively rare cases in which the security failure is critical enough to cause huge losses. In our model we describe the situation in terms of its extreme behavior. For instance, a possible risk situation can be characterized as a mean daily loss of  $X$  dollars, while in our approach it is characterized as having a probability of  $Z\%$  for the daily loss to exceed  $Y$  dollars. Observations that first appear to be outliers may not in fact be inconsistent with the rest of the data if they come from a long-tailed distribution. In some decisions it is the extreme values that

are of primary interest, as in the case of security failures, which are low-probability events that can bring huge losses.

We apply the notion of VaR and extreme value theory to security investment decisions. VaR is widely used in financial engineering and insurance (Crouhy et al. 2001, Duffie and Pan 1997, Jorion 1997, Dowd 1998, Embrechts 1996, de Fontnouvelle et al. 2005, Holton 2003). It has also been applied in airline risk management (Hallerbach and Menkveld 1999) and agriculture (Manfredo and Leuthold 1998). Mitra and Wang (2005) employed VaR for network revenue management through stochastic traffic engineering. Extreme value theory is an important statistical area in applied sciences and has found applications in engineering (Castillo 1988), insurance and finance (Embrechts et al. 1997), and management strategy (Dahan and Mendelson 2001), as well as in environmental and biomedical research. We utilize VaR to measure the information security risk faced by a firm and extreme value analysis to estimate the VaR of daily losses.

## 2.2. Value at Risk

Table 1 lists the notations that we will use in the following discussion. Let  $j$  denote the type of incident,  $j = 1, 2, \dots, T$ , and  $n_j$  denote the daily number of incidents of type  $j$ .  $X$  denotes the daily number of activities, and  $P_j(I)$  denotes the occurrence probability of an activity resulting in an incident of type  $j$  given a security investment level  $I$ . We have  $n_j = X \cdot P_j(I)$ . Security measures, such as firewalls, intrusion detection systems, multifactor authentication, multi-player defense, etc., reduce the occurrence probability  $P_j(I)$ , thus reducing  $n_j$ . Let  $C_j(I)$  denote the cost caused by an incident of type  $j$ . Through risk management practices such as disaster planning, data backup and recovery, redundancy/diversity, etc., we try to limit the impact of a security incident. Many studies have been carried out to estimate the occurrence probability  $P_j(I)$  and the incident impact  $C_j(I)$ : Kesh et al. (2002) developed a framework for analyzing e-commerce security by examining the relationships among e-commerce security needs, threats, technologies, and tools. Geer et al. (2003) introduced business-adjusted risk (BAR) for classifying security defects by

**Table 1** Notations

Value at risk	
$j$	The type of incident, $j = 1, 2, \dots, T$ .
$n_j$	The daily number of incidents of type $j$ .
$X$	The daily number of activities.
$I$	The security investment level.
$P_j(I)$	The occurrence probability of an activity resulting in an incident of type $j$ at security investment level $I$ .
$C_j(I)$	The cost caused by a successful exploit of type $j$ at security investment level $I$ .
$L$	The daily loss.
$F_L(z)$	The cumulative distribution function (CDF) of $L$ .
$z_p$	The $p$ th quantile of $F_L(z)$ . $p$ is a probability, such that $0 \leq p \leq 1$ .
Extreme value analysis	
$t$	Day $t$ , $t = 1, 2, \dots, 430$ .
$x_t$	$x_t$ is the observed daily loss in day $t$ , $t = 1, 2, \dots, 430$ .
$\Delta x_t$	$= x_t - x_{t-1}$ , $t = 2, \dots, 430$ .
$u$	High threshold.
GPD	Generalized Pareto distribution.
$\sigma$	Scale parameter of GPD.
$\xi$	Shape parameter of GPD.

their vulnerability type, degree of risk, and potential business impact. Farahmand et al. (2005) presented a subjective analysis, probability assessment, and damage evaluation of information security incidents. Sun et al. (2006) developed an evidential reasoning approach under the Dempster-Shafer theory of belief functions for information system risk assessment.  $P_j(I)$  and  $C_j(I)$  may be characterized as random variables following certain distributions. We consider  $P_j(I)$  and  $C_j(I)$  to be endogenous variables that can be adjusted through our security investment.

The daily loss includes all losses caused by different incidents/exploits in one day. Given a level of security investment  $I$ , the daily loss  $L$  is then

$$L = \sum_{j=1}^T n_j C_j(I) = X \sum_{j=1}^T P_j(I) \cdot C_j(I). \quad (1)$$

We assume that the incidents/exploits of different types are independent.

With proper security investment, we expect that the daily loss above a given level will be at a desired probability. Denote the cumulative distribution function (CDF) of  $L$  by  $F_L(z)$ . We define the VaR of the daily loss with probability  $p$  as

$$p = \Pr[L \geq VaR] = 1 - \Pr[L \leq VaR] = 1 - F_L(VaR) \quad (2)$$

From the definition, the probability that the organization will encounter a daily loss that exceeds VaR is  $p$ . Alternatively, we can say that with probability  $(1-p)$ , the daily loss encountered by the organization is less than or equal to VaR. The definition shows that VaR is concerned with tail behavior of the CDF  $F_L(z)$ . More specifically, the right tail of  $F_L(z)$  is our concern.

Given a univariate CDF  $F_L(z)$  and probability  $p$  ( $0 \leq p \leq 1$ ), the  $p$ th quantile of  $F_L(z)$  is defined as

$$z_p = \inf\{z \mid F_L(z) \geq p\}, \quad (3)$$

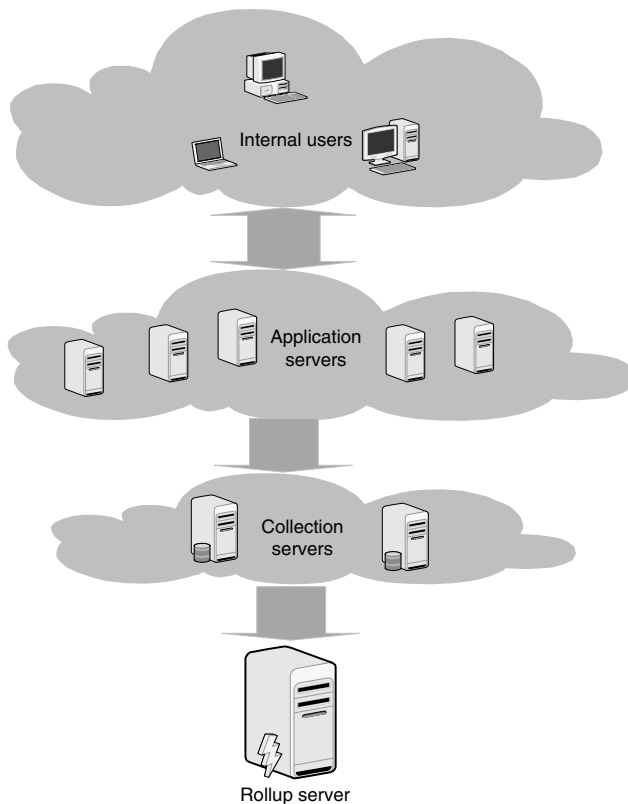
where  $\inf$  denotes the smallest real number satisfying  $F_L(z) \geq p$ . If the CDF  $F_L(z)$  is known, then the VaR is simply its  $(1-p)$ th quantile. However, we do not know the CDF in practice. Thus, study of the VaR of the daily loss is essentially concerned with estimation of the CDF  $F_L(z)$  and its quantile, especially the tail behavior of the CDF.

In this study we will use extreme value theory to characterize the stochastic behavior of daily losses and estimate their VaR. The decision maker can use this information to make a proper investment choice considering his/her risk preference.

### 3. Data

We collected data on daily activities between January 16, 2004, and March 20, 2005, from a large financial institute situated in the northeastern United States. These activities were recorded by a host-based activity monitoring system that monitors a set of internal servers including printer servers, application servers, file servers, and database servers. Figure 1 shows the architecture of the monitoring system. The institute has approximately 150 application servers that host different objects, including user objects, workstation objects, network components, and server components and subcomponents. Each application server has a monitoring agent running that collects user activity logs and sends them to a set of collection servers. Then a rollup server aggregates all logs from the collection servers and stores them in a large SQL database. An entity-relationship (ER) diagram of the database is shown in Figure 2. Recorded activities include the following:

- employee login/logout, file/printer access, or any other activity done at network level,

**Figure 1** The Architecture of the Activity Monitoring System

- inter-server communication (most of which happens automatically or is scheduled), and
- application access information (only some applications are monitored).

There are a total of 15,000 users of this system, distributed over two cities. Figure 3 plots the number of activities captured on a daily basis over 430 days. The mean of these daily numbers is 1,131,246.

In order to examine the abnormal activities and exploits, a set of daily reports were generated based on activity signatures. Table 2 shows an example of an SQL query used to create a daily report of password changing attempts. Appendix A lists four separate examples of such reports. On a daily basis, approximately 20 such daily monitoring reports covering a variety of activities are generated. Security analysts study these reports. Depending on business procedures (or decision rules), the analysts filter out false alarms and consequently identify the critical incidents. Those business procedures and or decision rules are developed and adjusted based on

business requirements. Critical incidents are reported to department managers for proper actions.

Based on our interviews with the institute's managers of information security, we classify the incidents into two types:

- Type A: High-Frequency–Low-Impact incidents. This includes incidents/exploits that affect individual users, such as changes to user objects. An incident of this type is expected once every 15,000 activities; on average, there are approximately 75 such incidents per day. The loss resulting from incidents of this type includes the cost of help-desk calls and loss of individual productivity and is estimated to be \$100 per incident on average.

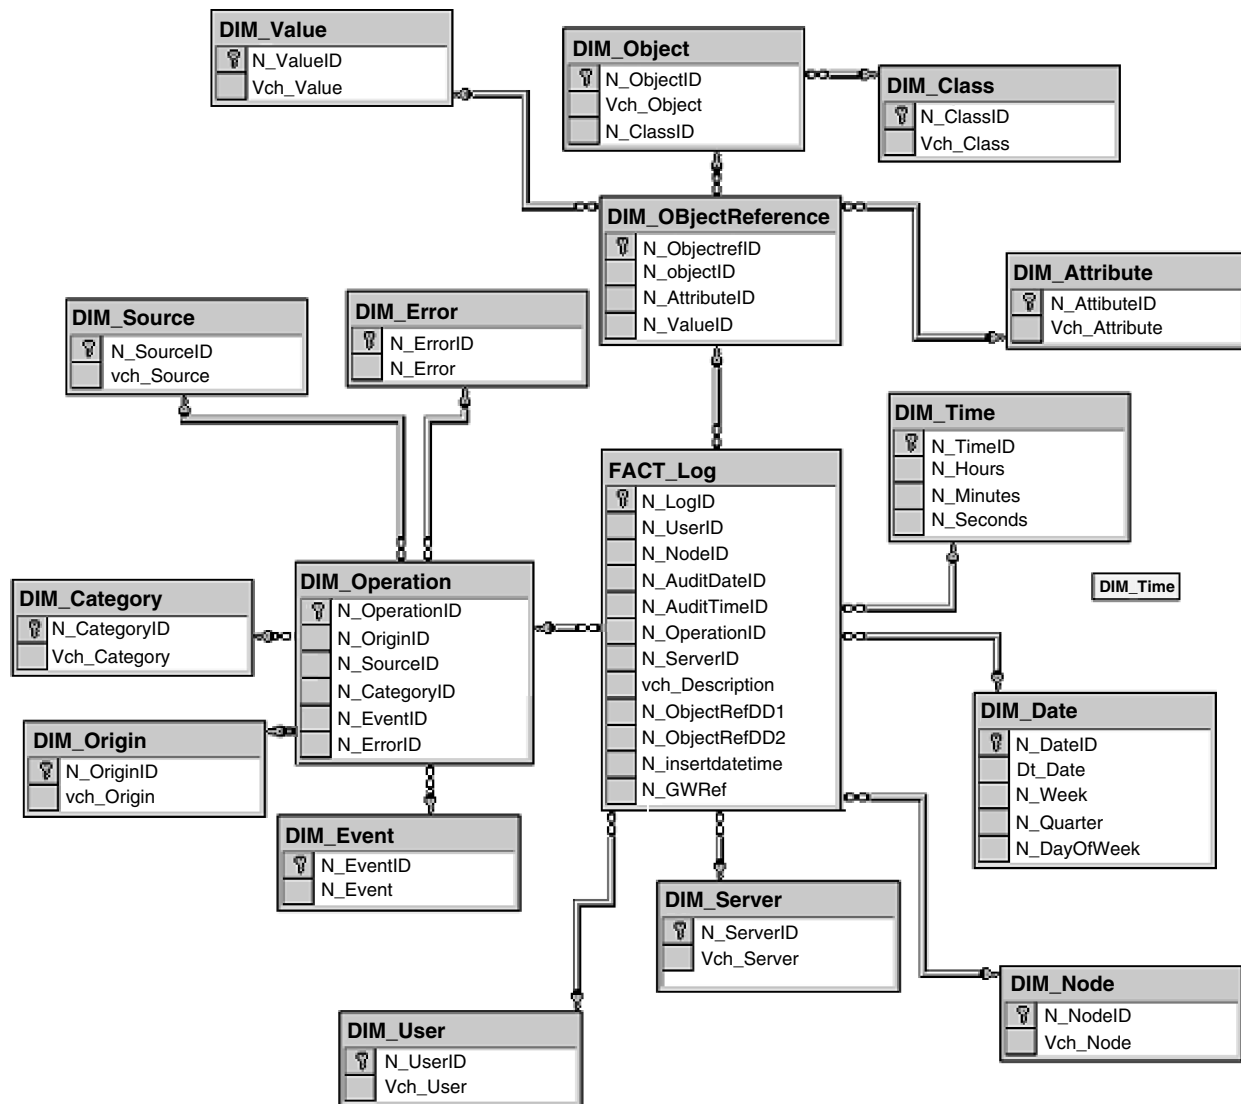
- Type B: Low-Frequency–High-Impact incidents. This includes incidents/exploits that cause disruptions of network services, such as e-mail server outages or network connection interruptions. These incidents affect a group of users. An incident of this type is expected once every 50,000,000 activities; on average, there will be one or two such incidents every two months. On an average the cost of this type of incident is approximately \$10,000 per incident, which includes the loss of user productivity, business disruptions, and possible damage to business reputation.

Based on this information and the daily number of activities, we simulated the daily loss and plotted it in Figure 4. The mean of the simulated daily losses is \$7,777. Further analyses are carried out below with this set of data.

#### 4. Extreme Value Analysis

In this section we will use extreme value theory to characterize the tail behavior (extremely large) of the daily losses. We first apply the augmented Dickey-Fuller test (Dickey and Fuller 1979, 1981) to examine the stationarity of the series (the daily losses). Then, the exceedances over thresholds extreme value model (Davison and Smith 1990, Pickands 1975) is fitted to the data, both with and without the assumption of the temporal independence of the observations. Seasonal phenomena (weekend and weekday) and time trends are examined. Based on our interviews with security managers at the financial institution we use a simple investment example to illustrate the method. Sensitivity analysis is also performed.

Figure 2 Entity-Relationship Diagram of the Activity Monitoring System Database



#### 4.1. The Augmented Dickey–Fuller Test


Because the model selection and the model estimation depend on whether the time series is stationary or not, we will begin by using the augmented Dickey–Fuller test to explore the data.

A univariate time series is said to be stationary if its mean, variance, and auto-covariance are independent of time. The first two conditions require the process to have a constant mean and variance, respectively, whereas the third one requires the covariance between any two values to depend only on the time interval between these two values and not on the


point in time  $t$ . A nonstationary time series does not satisfy one or all of these conditions; it has characteristics that change systemically through time. This nonstationarity may be caused by seasonal effects and/or long-term trends. If the series is not stationary, we will difference it or incorporate other possible factors into the modeling.

The augmented Dickey–Fuller (ADF) test examines whether a series is stationary (Greene 2000). It is a modification of the Dickey–Fuller test (Dickey and Fuller 1979) and involves lagged values of the dependent variable. The test captures the possibility that

**Figure 3** Daily Activities from 1/16/04 to 3/20/05



**Figure 4** Simulated Daily Losses



the observations are characterized by a higher-order autoregressive process. Equation (4) specifies what we used to compute an ADF:

$$\Delta x_t = \mu + \gamma x_{t-1} + \beta t + \sum_{j=1}^k \delta_j \Delta x_{t-j} + \varepsilon_t \quad (4)$$

where  $x_t$  is the observed daily loss at day  $t$  and  $\Delta x_t = x_t - x_{t-1}$ .  $\mu$  is a constant term.  $\gamma$ ,  $\beta$ , and  $\delta_j$  ( $j = 1, 2, \dots, k$ ,  $k$  is the number of lag) are coefficients.  $\varepsilon_t$  is the error term. The null hypothesis of the test is  $\gamma = 0$ , which implies that the data are not stationary, and the alternative hypothesis is  $\gamma < 0$ , which says that the series is stationary.

The trend and constant terms were not included in our test. Several values for  $k$  were used to examine whether the test results were affected by the number

of lags. From the results summarized in Table 3, we found that the series of the daily losses is stationary.

**4.2. Exceedances over Thresholds Model**

Because a series of daily data is available, we will employ the model known as “exceedances over thresholds” (Davison and Smith 1990, Pickands 1975) to fit the data. (In Appendix B, we provide a brief review of the exceedances over thresholds model.) The maximum-likelihood method is used to estimate the distribution parameters of generalized Pareto distribution (GPD) with the S-PLUS functions obtained from the website (Coles 2001).

**4.2.1. With the Assumption of Temporal Independence.** To use the exceedances over thresholds model, a proper threshold must be selected. The mean residual life plot is a diagnostic plot drawn before fitting any model that gives guidance about what threshold to use (Coles 2001, p. 78). Above a threshold  $\mu_0$  at which the generalized Pareto distribution provides a valid approximation to the excess distribution, the mean residual life plot should be approximately linear in the threshold chosen  $\mu$ . Figure 5 shows the

**Table 2** An SQL Query to Generate a Daily Report on Password Change Exceptions

---

Report Query Description: Password Changes Report  
 Output To: Screen  
 Report Query Selection String: SELECT LogDate AS ReportDate, LogTime as ReportTime, Origin, Source, Category, Event, Error, [User], Node, Server, substring([Remarks], 1, 255) AS Description1, substring([Remarks], 256, 510) AS Description2, substring([Remarks], 511, 765) AS Description3, substring([Remarks], 766, 1000) AS Description4, Object1, Class1, Attribute1, Value1, Object2, Class2, Attribute2, Operation, PerformedOn  
 FROM VIEW\_LTA\_REPORT\_80 WHERE ((Origin LIKE 'LT Auditor+ Processed Log' AND Source LIKE 'NetWare' AND Category LIKE 'NDS' AND Event = 409)) AND (LogDate >= '2005-02-27 12:00:00 AM' AND LogDate <= '2005-02-27 11:59:59 PM')


---

**Table 3** Results of the Augmented Dickey-Fuller Test

$k$	1	5	10
ADF $t$ statistics	-9.88	-4.45	-2.48

*Notes.* The critical value at 1% is -2.57. The critical value at 5% is -1.94. The critical value at 10% is -1.61.

Figure 5 Mean Residual Life Plot for Daily Losses



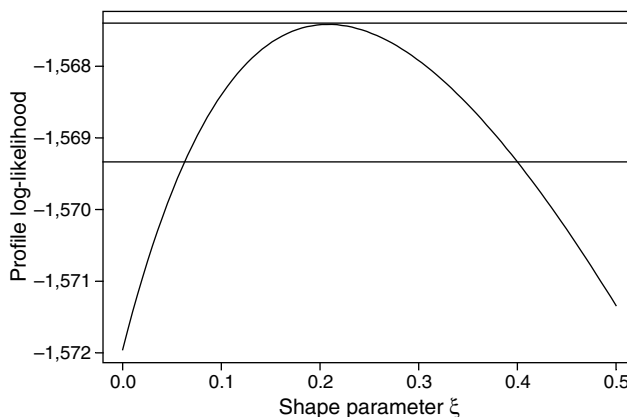
mean residual life plot with approximate 95% confidence intervals for daily losses.

The graph is initially linear from  $\mu = 0$  to  $\mu \approx 8,500$  but shows substantial curvature until  $\mu \approx 24,500$ , whereupon it decays sharply. It is tempting to conclude that there is no stability until  $\mu = 24,500$ , after which there is approximate linearity. However, there are only two exceedances of the threshold  $\mu = 24,500$ , too few to make meaningful inferences. Furthermore, the information in the plot for large values of  $\mu$  is unreliable because of the limited amount of data on which the estimate and confidence interval are based. Thus, we chose  $\mu = 8,500$  as our threshold. (We also did a sensitivity analysis with  $\mu = 8,700$ ; the results do not have significant differences.) This choice leads to 178 exceedances in the series of length 430. The maximum likelihood estimators of GPD parameters are  $(\hat{\sigma}, \hat{\xi}) = (1,993.16, 0.21)$ , with standard error 223.83 and 0.09, respectively. Figure 6 plots the profile log-likelihood for  $\xi$  with a line indicating its 95% confidence interval.

Diagnostic plots for the fitted GPD are shown in Figure 7. The goodness-of-fit in the quantile plot seems unconvincing, but the confidence intervals on the return level plot suggest that the model departures are not large after allowance for sampling. The return level plot also illustrates the very large uncertainties that accrue once the model is extrapolated to higher levels.

The GPD model provides a direct method for risk estimation using return level. Figure 8 plots the profile log-likelihood for the 100-day return level  $\hat{x}_{0.99}$  with a

Figure 6 Profile Likelihood for Shape Parameter in Threshold Excess Model of Daily Losses

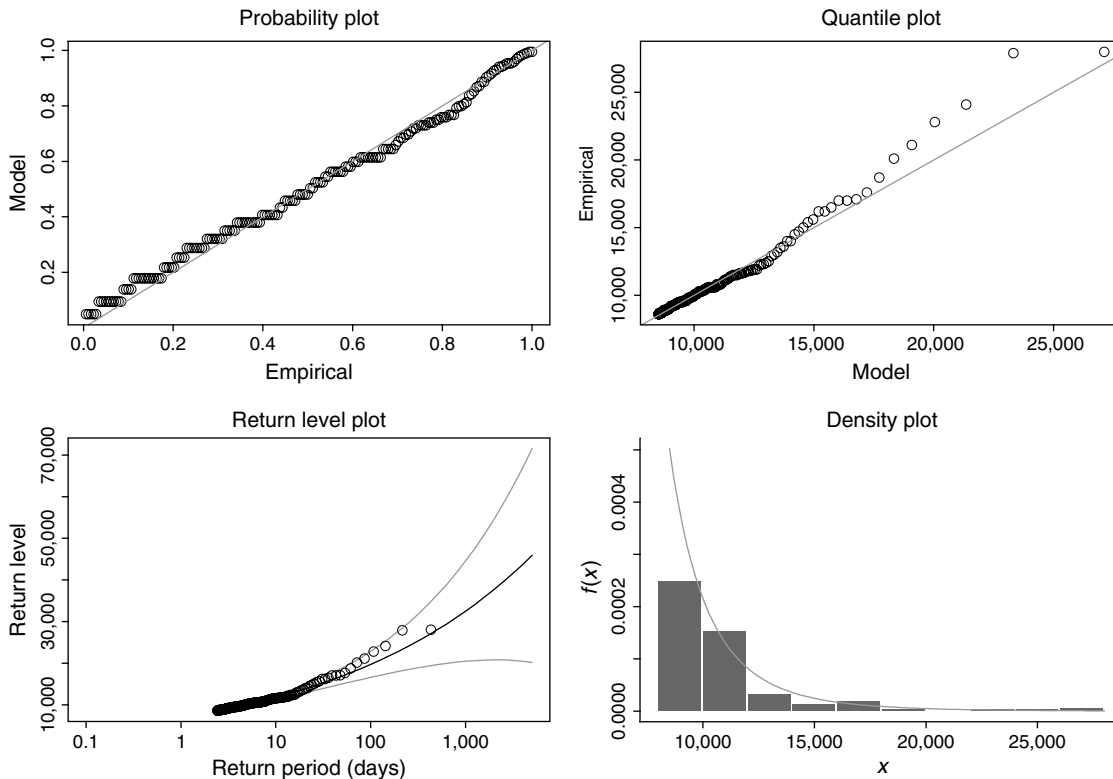


line indicating its 95% confidence interval. In our analysis we use 100 days as our return period. In other words, we estimate a level at which the occurrence probability of the daily losses is greater than or equal to 1%. We have  $\hat{x}_{0.99} \approx \$19,710$  with standard error \$2,692. This is the VaR of daily loss at  $p = 1\%$ .

**4.2.2. Without the Assumption of Temporal Independence.** In the above analysis we assume that our observations meet the property of temporal independence. However, this may not be the case in reality, and the observations may be mutually dependent. Various methods with differing degrees of sophistication have been suggested to deal with the problem of dependent exceedances in the threshold model. The most widely used method is declustering (Coles 2001, Davison and Smith 1990). The purpose of declustering is to filter the dependent observations and obtain a set of threshold excess that are approximately independent. The cluster may be identified through a parametric model that utilizes the extremal index (Davison and Smith 1990, Leadbetter et al. 1989, Smith and Weissman 1994) or a nonparametric method in which the clusters of exceedances are defined by an empirical rule and the extremal index is then estimated as the reciprocal of mean cluster size. Several studies have shown that the nonparametric method is quite robust, provided that the empirical rules are intuitively reasonable (Smith 1989, Tawn 1988).

In our analysis we will follow the nonparametric method. First we specify a threshold  $u$  and define consecutive exceedances of  $u$  to belong to the same

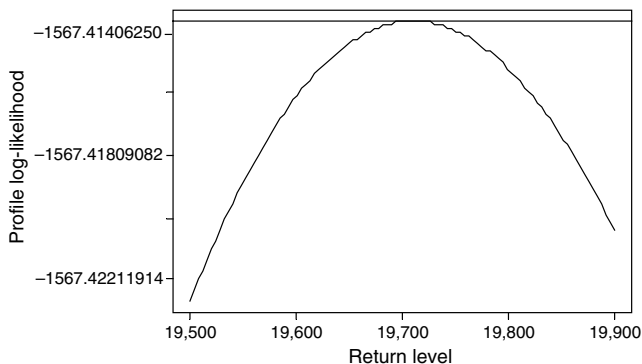
Figure 7 Model Diagnostic Plots for Threshold Excess Model Fitted to Daily Losses



cluster. Once we obtain  $r$  consecutive observations that fall below  $u$ , the cluster is terminated. The next exceedance of  $u$  will start a new cluster. After we obtain clusters, we identify the maximum value within each cluster. Assuming the cluster maxima to be independent, we fit the generalized Pareto distribution to the cluster maxima.

We then need to choose two values for the nonparametric method of declustering. To check the sensitivity

Figure 8 Profile Likelihood for 100-Day Return Level in the Threshold Excess Model



of the estimate, we did a full-factor analysis for  $u = 8,500, 8,700$ , and  $r = 0, 2$ . Because of the restriction of sample size, we were unable to make  $r$  larger to obtain a meaningful estimation. In the case of  $r = 0$ , each exceedance forms a cluster, which corresponds to the analysis we carried out in §4.2.1.

Table 4 shows the estimation results, including the standard error (SE). The values for 100-day return level ( $\hat{x}_{0.99}$ ) are similar across the choices of  $u$ . The values for 100-day return level ( $\hat{x}_{0.99}$ ) at  $r = 2$  are lower than that at  $r = 0$ , but the differences are not statistically significant. This suggests that inference on return levels is quite robust despite the subjective choice that needs to be made.

**4.2.3. Seasonal Phenomena and Time Trends.** To examine whether the exceedances over the threshold vary between weekends and weekdays (seasonal phenomena) and/or change linearly over time, we will model the GPD parameters in generalized linear models. Let  $y_t$  be exceedances over the threshold, i.e.,  $y_t = x_t - u$ , where  $x_t$  is the daily loss at day  $t$  and  $u$  is

**Table 4 Estimation Results**

	$u = 8,500$		$u = 8,700$	
	$r = 0$	$r = 2$	$r = 0$	$r = 2$
The number of clusters ( $n_c$ )	178	56	163	53
GPD parameter ( $\hat{\sigma}$ ) (SE)	1,993.16 (223.83)	3,180.23 (674.29)	1,954.16 (233.60)	3,309.96 (711.20)
GPD parameter ( $\hat{\xi}$ ) (SE)	0.21 (0.09)	0.17 (0.17)	0.23 (0.09)	0.15 (0.17)
100-day return level ( $\hat{x}_{0.99}$ ) (SE)	19,710.23 (2,692.80)	13,582.85 (3,391.10)	19,813.89 (2,642.37)	13,818.25 (3,810.25)

the threshold. Define a dummy variable  $I(t)$  as

$$I(t) = \begin{cases} 1, & \text{if } t \text{ is Saturday or Sunday} \\ 0, & \text{otherwise} \end{cases}$$

Then the estimation model is

$$y_t \sim \text{GPD}(\sigma(t), \xi(t))$$

where

$$\sigma(t) = \beta_{10} + \beta_{11}I + \beta_{12}t$$

$$\xi(t) = \beta_{20} + \beta_{21}I + \beta_{22}t.$$

The estimators of the coefficients  $\beta_{ij}$  ( $i = 1, 2$ , and  $j = 0, 1, 2$ ) and their standard error are shown in Table 5. The estimation is carried out with  $u = 8,500$ . We can see that none of the estimators for  $\beta_{ij}$  ( $i = 1, 2$ , and  $j = 1, 2$ ) are significant at 1%. This implies that the exceedances over the threshold neither vary between weekdays and weekends (seasonal phenomena) nor change linearly over time. The results are also consistent with the result of the augmented Dickey-Fuller test.

### 4.3. A Simple Investment Example

Consider two possible security investment solutions:

- Solution A: Install a firewall. This solution will decrease the probability of both types of incidents by half at a daily cost of \$550.

**Table 5 Estimated Coefficients of the Linear Models for GPD Parameters**


$ij$	10	11	12	20	21	22
$\hat{\beta}_{ij}$	1,936.82	-269.76	1.06	-0.26	0.17	0.00
(SE)	(643.80)	(487.95)	(2.53)	(0.21)	0.18	(0.00)

**Table 6 Solutions A and B**

	VaR (at $p = 1\%$ ) of daily loss (\$)	Expected daily cost (\$)*
Current	19,710	7,777
Solution A	11,876	4,475 (=3,925 + 550)
Solution B	6,938	4,906 (=2,906 + 2,000)

\*Expected daily cost = Average daily loss + Solution daily cost.

**Figure 9 Solution Plot**



- Solution B: Increase the frequency of backup and introduce detection systems. This solution will decrease losses from both incident types by half and reduce the frequency of type A intrusions to 1 per 20,000 activities. It has a daily cost of \$2,000. Following the procedure in §4.2.1, we estimated the VaR and calculated the expected daily cost (which includes the average daily loss and daily solution cost) for both solutions. The results are shown in Table 6.

The current status and both solutions are plotted in Figure 9. We can see that, compared to the current status, both solutions reduced expected daily cost as well as value at risk. Solution A has a higher VaR and a lower expected daily cost than solution B. The decision maker can choose either A or B based on his/her risk preference, but if the decision maker follows the rule of minimizing expected daily cost, only solution A is suggested.

### 4.4. Sensitivity Analysis

To examine the sensitivity of the result, we further interviewed the security managers about the occurrence probability and incident loss. Based on the

**Table 7** Parameters for Sensitivity Analysis

Incident type	Occurrence probability	Incident loss
Type A: High-frequency–low-impact incidents	Triangular distribution with mode 1/15,000, minimum 1/20,000, and maximum 1/12,000.	Uniform distribution with minimum \$90 and maximum \$110.
Type B: Low-frequency–high-impact incidents	Normal distribution with mean 1/50,000,000 and standard deviation 1/500,000,000.	Normal distribution with mean \$10,000 and standard deviation \$1,000.

interviews and data snapshots, we used a set of random variables as shown in Table 7 to represent the occurrence probability and incident loss. Daily losses were then simulated.

The series remains stationary. We performed the analysis following the steps in §4.2.1. The expected daily loss in current status is \$7,752, and the VaR at  $p = 1\%$  (or the 100-day return level) of daily losses is \$20,175 (with a standard deviation of \$3,842).

Consider that solution A will reduce the occurrence probability of both types of incidents by half (but not change the distribution of the probability) at a daily cost of \$550. Solution B will decrease the loss from both types by half and reduce the frequency of type A intrusions to a triangular distribution with mode 1/20,000 activities, minimum 1/25,000, and maximum 1/17,000, with a daily cost of \$2,000.

With solution A, the expected daily loss is \$3,908, and the total expected daily cost is \$4,458. The VaR of daily loss at  $p = 1\%$  is \$10,304 with a standard deviation of \$1,274. With solution B, the expected daily loss is \$2,904, and the total expected daily cost is \$4,904. The VaR of daily loss at  $p = 1\%$  is \$7,788 with a standard deviation of \$1,119. These results are similar to the ones calculated in §4.3.

## 5. Discussion and Conclusion

Learning from the past gives us a way to prepare for the future. Extreme value analysis helps us to quantify the risk of information security. Using this quantification, a firm can determine proper security solutions based on its risk preference.

This methodology can lead to many future avenues of research. With the extreme value approach, we may be able to determine whether the extreme daily loss is influenced by other environmental factors,

which would help to make strategic investment in information security more effective. Another possible extension is to apply the extreme value analysis to intrusion detection, another important area of information security. Detecting and understanding anomalies on the Internet is still an open and ill-defined question.

This study has certain limitations. First, for illustration purposes, we have used simulated daily losses in our analysis. Those simulated daily losses may not reflect the real losses to the institutions. Actual losses in a firm will be situation specific and can be estimated only on the basis of assumptions by security staff, many of which would not be verifiable. However, an indicative figure as used in our simulation can provide insight facilitating managerial judgment. Second, we focused on using extreme value theory to characterize the behavior of attacks; we did not examine how to operationally decide on a corresponding security investment level and then convert it into a real protection level for a system through the combination of various technology and security policies. No single security technology can achieve effective security; therefore, diverse security technologies that substitute for and complement each other (e.g., firewalls, intrusion detection systems, monitoring technologies) are necessary. The ability of such technologies to reduce system risk depends on many factors that are not considered in the model, such as how the technologies collaborate with other security mechanisms. This is an interesting topic that needs further exploration. Third, because extreme value analysis is a statistical approach based on past data, it has limited applications in cases where the security scenario is evolving so that past data are no longer reliable predictors of future situations.

## Acknowledgments

An earlier version of this paper was presented at the International Conference on Information Systems, December 11–14, 2005, Las Vegas, NV. The authors would like to thank the session chair, Professor F. Farahmand, and the discussant, Professor Q. Hu, for their comments. The authors also thank Professors N. Menon, K. Dogan, Y. Ryu, J. Asundi, W. Yue, S. Sarkar, S. Menon, H. Cavusoglu, S. Raghunathan, R. Mookerjee, and V. Jacob, as well as other participants in a seminar at the University of Texas at Dallas, for several insightful comments that helped improve the paper. Last but not least, the authors thank Mr. M. Gupta,

systems security manager at a large financial institution in the northeast United States, for various discussions that gave us insight into the security aspects that we examined in this paper. This research has been funded in part by the National Science Foundation under Grants 0402388 and 0705292. The usual disclaimer applies. Most importantly the authors thank the senior editor, the associate editor, and the three referees for their comments, which have considerably improved the lucidity of this research.

### Appendix A. Examples of Security Incident Reports

We contacted the financial institution for permission to see their exception reports. After consulting with the chief information security officer, our contact at the financial institution gave us specific reports for a sample of five randomly chosen days, which we have summarized in Table A.1 below and have described in the subsequent paragraphs.

#### 1. Web Banking—Bill Pay Files Monitoring

Bill pay files are used by commercial banks and institutions to verify the charges occurring on their customers' debit (or credit) cards. These files contain customers' transaction information. Business entities such as retail websites send these transaction files to the institution periodically (e.g., every day).

These files are stored in a local server belonging to the commercial bank or institution. Normally the verification process is executed at night. These files should not be viewed or changed by anyone. Viewing these files results in information leakage. Modification or deletion of these files may cause financial loss to the banks. To monitor any unauthorized actions (such as view, modification, or deletion) on these files, daily reports are generated. The daily numbers of such exceptions in five observation days are 1, 1, 1, 0, and 2, respectively. All of those incidents had low impact.

#### 2. CORPX Activities Report

CORPX is a technical firm who provides service to the institutions. The institutions have outsourced part of their helpdesk to CORPX. CORPX can access the network infrastructure of the institutions. Based on the regulations and laws, the activities of CORPX inside of the institution infrastructure should be monitored, and any exception or unauthorized activities should be reported. Daily reports are

generated to monitor those exceptions on the activities of CORPX inside the infrastructure. The daily numbers of such exceptions in five observation days are 14, 8, 1, 15, and 5, respectively. All of those incidents had low impact.

#### 3. NDS Security Equivalence Exception

NDS (network directory service) security equivalence exception monitors the activities of security staffs as they manage critical roles (assign and revoke) in a role-based access control system. There are approximately 300 such critical roles, which have access to important resources. To prevent abuse of access rights, assigning or revoking these roles needs to be monitored. The daily numbers of such exceptions in five observation days are 5, 12, 3, 13, and 13, respectively. All of those incidents had low impact.

#### 4. NDS Object Changes

NDS (network directory service) provides an organizational-level structure of the enterprise information infrastructure. The objects in the NDS include servers, workstations, user/role objects, group objects, and network objects. Unauthorized changing or moving any object in NDS affects the normal function of the corresponding object. Depending on the nature of the objects, the business impact may be high (e.g., a server) or low (e.g., a user object). The report on NDS object changes monitors all exceptions relating to changes of these objects. The daily numbers of such exceptions in five observing days are 0, 1, 10, 28, and 2, respectively. On the third observation day, the SAS service on a server was accidentally deleted. The deletion interrupted the production activities of the customer analytics team. The team relies on SAS for marketing and risk analysis. It took around half a day to restore the SAS service. This exception was considered as a high-impact incident.

Table A.1 summarizes the daily numbers of those four types of exceptions.

### Appendix B

#### 1. Exceedances over Thresholds

Classical extreme value theory concerns block maxima, which considers only one observation in a sequence of observations (a block) (Coles 2001, Fisher and Tippett 1928, Gumbel 1958). Exceedances over thresholds provide an alternative way to model extreme value by characterizing an observation as extreme if it exceeds a high threshold. The concept of exceedance over thresholds considers an observation as extremal when it is above a certain threshold. If an entire time series of, say, hourly or daily observations is available, exceedances over thresholds provides a better utilization of the data. The major theorem regarding exceedances over thresholds is as follows:

**THEOREM (COLES 2001, P. 75).** *Consider the distribution of  $X$  conditionally on exceeding some high threshold  $u$ , and let  $Y = X - u$ , and  $Y > 0$ . We know*


$$F_u(y) = \Pr\{Y \leq y \mid Y > 0\} = \frac{F(u + y) - F(u)}{1 - F(u)}$$

**Table A.1** Daily Numbers of Four Types of Exceptions

	Web banking—Bill pay files monitoring	CORPX activities report	NDS security equivalence exception	NDS object changes
Day 1	1	14	5	0
Day 2	1	8	12	1
Day 3	1	1	3	10*
Day 4	0	15	13	28
Day 5	2	5	13	2

\*A high-impact incident happened.

**Figure B.1** Densities for the Three Generalized Pareto Distributions ( $\sigma = 1, u = 0$ )



As  $u \rightarrow \omega_F = \sup\{x: F(x) < 1\}$ , we found a limit distribution,

$$F_u(y) \approx G(y; \sigma, \xi)$$

where  $G$  is Generalized Pareto Distribution (GPD)

$$G(y; \sigma, \xi) = 1 - \left(1 + \xi \frac{y}{\sigma}\right)^{-1/\xi} \tag{B1}$$

defined on  $\{y: y > 0 \text{ and } (1 + \xi y/\sigma) > 0\}$ , where  $\sigma$  and  $\xi$  are the two parameters of the distribution.

The rigorous connection between exceedances over thresholds and the classical extreme value theory was established by Pickands (1975). GPD has three cases depending on the value of the parameter  $\xi$  (Figure B.1):

- The case  $\xi > 0$  is the “long-tailed” case, for which  $1 - G(x)$  decays at the same rate as  $x^{-1/\xi}$  for large  $x$ . This is reminiscent of the usual Pareto distribution,  $G(x) = 1 - cx^{-1/\xi}$ .
- For  $\xi = 0$ , we have the exponential distribution with mean  $\sigma$  as the limit

$$G(y; \sigma, 0) = 1 - \exp\left(-\frac{y}{\sigma}\right).$$

- For  $\xi < 0$ , the distribution has finite upper endpoint at  $-\sigma/\xi$ .

Replacing  $Y = X - u$  into Equation (B1), now we have

$$\Pr\{X > x \mid X > u\} = \left(1 + \xi \frac{x - u}{\sigma}\right)^{-1/\xi} \tag{B2}$$

It follows that

$$\Pr\{X > x\} = s_u \left(1 + \xi \frac{x - u}{\sigma}\right)^{-1/\xi} \tag{B3}$$

where  $s_u = \Pr\{X > u\}$ . By inverting Equation (B3) we obtain

$$x_{1-p} = \begin{cases} u + \frac{\sigma}{\xi} \left[ \left(\frac{s_u}{p}\right)^\xi - 1 \right], & \text{for } \xi \neq 0 \\ u + \sigma \log \frac{s_u}{p}, & \text{for } \xi = 0. \end{cases} \tag{B4}$$

$x_{1-p}$  is the  $(1/p)$ -observation return level. In other words, the level  $x_{1-p}$  is expected to be exceeded once every  $1/p$  observations to a reasonable degree of accuracy, or the occurrence probability of an observation to exceed  $x_{1-p}$  is  $p$ .

### 2. Seasonal Phenomena and Trend Analysis

In a lot of cases, the stochastic process may have characteristics that change through time. There may be a seasonal phenomenon or time trend present in the data. Or we may be interested in whether the exceedances over the threshold will be influenced by the changes of other environmental factors.

Let  $GPD(\sigma, \xi)$  denote the GPD distribution with parameters  $\sigma$  and  $\xi$ . Let  $u(t)$  denote threshold at time  $t$ , and let  $y_t$  denote the exceedance over threshold at time  $t$ ; i.e.,  $y_t = x_t - u(t)$ . To examine whether there is time trend of  $\sigma(t)$ , the estimation model is

$$y_t \sim GPD(\sigma(t), \xi)$$

where

$$\sigma(t) = \beta_0 + \beta_1 t.$$

To examine the presence of the seasonal phenomena of  $\sigma(t)$  with  $k$  seasons  $s_1, s_2, \dots, s_k$ , we may use the form

$$\sigma(t) = [I_1(t), I_2(t), \dots, I_k(t)] \begin{bmatrix} \beta_1 \\ \beta_2 \\ \dots \\ \beta_k \end{bmatrix}$$

where  $I_j(t)$  is the dummy variable having

$$I_j(t) = \begin{cases} 1, & \text{if } s(t) = s_j \\ 0, & \text{otherwise} \end{cases}, \quad j = 1, \dots, k.$$

To identify factors that might have significant impact on the exceedances over the threshold, the estimation model in general form is

$$y_t \sim GPD(\sigma(t), \xi(t))$$

where we specify a general linear model for  $\sigma(t)$  and  $\xi(t)$  as

$$\sigma(t) = [1, z_{11}(t), \dots, z_{1n}(t)] \begin{bmatrix} \beta_{10} \\ \beta_{11} \\ \dots \\ \beta_{1n} \end{bmatrix}$$

and

$$\xi(t) = [1, z_{21}(t), \dots, z_{2m}(t)] \begin{bmatrix} \beta_{20} \\ \beta_{21} \\ \dots \\ \beta_{2m} \end{bmatrix}.$$

$z_{ij}(t)$  is a factor to be examined (e.g., the time, the number of employees, and the number of enterprise applications in a different time period).

Using these regression models, we are able to identify whether the exceedances over threshold are changing over time and/or whether there are seasonal phenomena or other important factors that significantly change the behavior of those extremal.

## References

- Anderson, R. H., R. Brackney. 2004. *Understanding the Insider Threat*. RAND Corporation, Santa Monica.
- Castillo, E. 1988. *Extreme Value Theory in Engineering*. Academic Press, San Diego.
- Chinchani, R., A. Iyer, H. Q. Ngo, S. Upadhyaya. 2005. Towards a theory of insider threat assessment. *The 2005 Internat. Conf. Dependable Systems and Networks (DSN'05)*, Yokohama, Japan. IEEE Computer Society, Washington, D. C. 108–117.
- Coles, S. 2001. *An Introduction to Statistical Modeling of Extreme Values*. Springer-Verlag, London.
- Crouhy, M., D. Galai, R. Mark. 2001. *Risk Management*. McGraw-Hill, New York.
- Dahan, E., H. Mendelson. 2001. An extreme-value model of concept testing. *Management Sci.* **47**(1) 102–116.
- Davison, A. C., R. L. Smith. 1990. Models for exceedances over high thresholds (with discussion). *J. Roy. Statist. Soc.* **52** 393–442.
- de Fontnouvelle, P., J. Jordan, E. Rosengren. 2005. Implication of alternative operational risk modeling techniques. NBER Working Paper 11103, National Bureau of Economic Research, Cambridge, MA. Available at <http://www.nber.org/papers/w11103>.
- Devaraj, S., R. Kohli. 2002. *The IT Payoff*. Prentice Hall, Upper Saddle River, NJ.
- Dickey, D., W. Fuller. 1979. Distribution of the estimators for autoregressive time series with a unit root. *J. Amer. Statist. Assoc.* **74** 427–431.
- Dickey, D., W. Fuller. 1981. Likelihood ratio tests for autoregressive time series with a unit root. *Econometrica* **49** 1057–1072.
- Dowd, K. 1998. *Beyond Value at Risk; The New Science of Risk Management*. John Wiley & Sons, New York.
- Duffie, D., J. Pan. 1997. An overview of value at risk. *J. Derivatives* **4**(3) 7–49.
- Embrechts, P. 1996. Actuarial versus financial pricing of insurance. Working paper, The Wharton School, Philadelphia. Available at <http://fic.wharton.upenn.edu/fic/papers/96/9617.pdf>.
- Embrechts, P., C. Kluppelberg, T. Mikosch. 1997. *Modeling Extremal Events for Insurance and Finance*. Springer, New York.
- Ernst & Young. 2003. Global Information Security Survey 2003, Ernst & Young LLP.
- Ernst & Young. 2004. Global Information Security Survey 2004, Ernst & Young LLP.
- Farahmand, F., S. B. Navathe, G. P. Sharp, P. H. Enslow. 2005. A management perspective on risk of security threats to information systems. *Inform. Tech. Management* **6**(2–3) 203–255.
- Fisher, R. A., L. H. C. Tippett. 1928. *Limiting Forms of the Frequency Distributions of the Largest or Smallest Member of a Sample*. The Cambridge Philosophical Society, Cambridge University Press, London.
- Gal-or, E., A. Ghose. 2005. The economic incentives for sharing security information. *Inform. Systems Res.* **16**(2) 186–208.
- Geer, D., K. S. Hoo, A. Jaquith. 2003. Information security: Why the future belongs to the quants. *IEEE Security & Privacy* **1** 32–40.
- Gordon, L. A., M. P. Loeb. 2002. The economics of information security investment. *ACM Trans. Inform. Systems Secur.* **5**(4) 438–457.
- Gordon, L. A., M. P. Loeb, W. Lucyshyn, R. Richardson. 2005. 2005 CSI/FBI Computer Crime and Security Survey. Computer Security Institute, San Francisco.
- Greene, W. H. 2000. *Econometric Analysis*. Prentice Hall, Upper Saddle River, NJ.
- Gumbel, E. J. 1958. *Statistics of Extremes*. Columbia University, New York.
- Hallerbach, W., B. Menkveld. 1999. Value at risk as a diagnostic tool for corporates: The airline industry. Papers No. 99-023/2, Tinbergen Institute Discussion Papers, Rotterdam, The Netherlands. Available at <http://www.tinbergen.nl/discussionpapers/99023.pdf>.
- Holton, G. A. 2003. *Value at Risk: Theory and Practice*. Academic Press, London.
- Hoo, K. J. S. 2000. How much is enough? A risk-management approach to computer security. Working paper, Center for International Security and Cooperation, Stanford University. Available at <http://iis-db.stanford.edu/pubs/11900/soohoo.pdf>.
- Jorion, P. 1997. *Value at Risk*. McGraw-Hill, New York.
- Kannan, K., R. Telang. 2005. Market for software vulnerabilities? Think again. *Management Sci.* **51**(5) 726–740.
- Kesh, S., S. Ramanujan, S. Nerur. 2002. A framework for analyzing e-commerce security. *Inform. Management Comput. Secur.* **10**(4) 149–158.
- Leadbetter, M. R., I. Weissman, L. De Haan, H. Rootzen. 1989. On clustering of high levels in statistically stationary series. *The 4th Internat. Meeting on Statist. Climatology*, New Zealand Meteorological Service, Wellington, New Zealand.
- Longstaff, T. A., C. Chittister, R. Pethia, Y. Y. Haimes. 2000. Are we forgetting the risks of information technology? *IEEE Comput.* **33**(12) 43–51.
- Manfredo, M. R., R. M. Leuthold. 1998. Agricultural applications of value-at-risk analysis: A perspective. *The NCR-134 Conf. Appl. Commodity Price Anal., Forecasting, and Market Risk Management*. St. Louis.
- Mercuri, R. T. 2003. Analyzing security costs. *Comm. ACM* **46**(6) 15–18.
- Mitra, D., Q. Wang. 2005. Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management. *IEEE/ACM Trans. Networking* **13**(2) 221–233.

- Pickands, J. 1975. Statistical inference using extreme order statistics. *Ann. Statist.* **3** 119–131.
- Schechter, S. E., M. D. Smith. 2003. How much security is enough to stop a thief? The economics of outsider theft via computer systems networks, *Proc. 7th Financial Cryptography Conf.*, Guadeloupe, French West Indies. 122–137.
- Shaw, E. D., K. G. Ruby, J. M. Post. 1998. The insider threat to information systems. *Security Awareness Bull.* **2-98**.
- Smith, R. L. 1989. Extreme value analysis of environmental time series: An example based on ozone data (with discussion). *Statist. Sci.* **4** 367–393.
- Smith, R. L., I. Weissman. 1994. Estimating the extremal index. *J. Roy. Statist. Soc. B*(56) 515–528.
- Sun, L., R. P. Srivastava, T. J. Mock. 2006. An information systems security risk assessment model under Dempster-Shafer theory of belief functions. *J. Management Inform. Systems* **22**(3) 190–142.
- Tawn, J. A. 1988. An extreme value theory model for dependent observations. *J. Hydrology* **101** 227–250.
- Varian, H. R. 2004. System reliability and free riding. L. J. Camp, S. Lewis, eds. *Economics of Information Security*. Kluwer Academic Publishers, Boston/Dordrecht/London, 1–15.