Computers & Operations Research 39 (2012) 774–784

Contents lists available at ScienceDirect

Computers & Operations Research

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

Quantifying information security risks using expert judgment elicitation

Julie J.C.H. Ryan [a], Thomas A. Mazzuchi [a,*], Daniel J. Ryan [b], Juliana Lopez de la Cruz [c], Roger Cooke [c,d]

[a] School of Engineering and Applied Science, The George Washington University, USA
[b] Information Resources Management College, National Defense University, USA
[c] Delft University of Technology, Netherlands
[d] Resources for the Future, USA

ARTICLE INFO

Available online 9 December 2010

Keywords:
Information security
Risk management
Probability distributions
Expert elicitation
Poisson processes

ABSTRACT

In the information security business, 30 years of practical and theoretical research has resulted in a fairly sophisticated appreciation for how to judge the qualitative level of risk faced by an enterprise. Based upon that understanding, there is a practical level of protection that a competent security manager can architect for a given enterprise. It would, of course, be better to use a quantitative approach to risk management, but, unfortunately, sufficient quantitative data that has been scientifically collected and analyzed does not exist. There have been many attempts to develop quantitative data using traditional quantitative methods, such as experiments, surveys, and observations, but there are significant weaknesses apparent in each approach. The research described in this paper was constructed to explore the utility of applying the well-established method of expert judgment elicitation to the field of information security. The instrument for eliciting the expert judgments was developed by two information security specialists and two expert judgment analysis specialists. The resultant instrument was validated using a small set of information security experts. The final instrument was used to elicit answers to both the calibration and judgment questions through structured interviews. The data was compiled and analyzed by a specialist in expert judgment analysis. This research illustrates the development of prior distributions for the parameters of models for cyber attacks and uses expert judgment results to develop the distributions.

1. Introduction

In their survey article on Operations Research models in Homeland Security, Wright et al. [1] identified the area of Cyber Security as a fruitful and important area for research under the general area of Critical Infrastructure Protection. The authors cite only a few articles in the literature that currently deal with operations research techniques for addressing this problem area, and of these, only Shindo et al. [2] approached a more formal risk analysis. Thus, formal risk analysis in the area of Cyber Security is a relatively new area for the operations research community, though risk analysis in other areas such as health risk [3], transport risk, [4], flight risk [5], programmatic risk [6], business process risk [7], and risk in product development [8] among others, have appeared in the operations research literature.

In the profession of information security, it has long been a goal to determine how to measure the value of investments in information security-related technology and practices. For example, Davis et al. [9], in an attempt to develop a security investment valuation process, examined the behavior of online customers in the wake of security incidents. For the purposes of their research, they limited the definition of security incident to that of a disclosed security breach falling into one of two types: "hack" and "fraud". The results of their analysis, which was based on the development of a behavior model through the application of Bayesian Markov Chain Monte Carlo sampling and subsequent analysis of the model properties with varying incident rates, reflected no return on investment proposition for enterprises. In their words: "our findings illustrate how difficult it is to convince corporations to invest in cyber security. The direct financial loss of revenue from lost customers seems to be not occurring." Their research, while interesting, is somewhat constrained. The concept of availability is one that is recognized in the security community as an important security attribute and yet was not considered to be such in this research. Further, their financial analysis was based on changes in stock prices of the corporation reporting the breach without controlling for any of the myriad of problems that might have been either supporting or distressing the trading value of the company. Thus while this research is interesting, it is ultimately of little use in understanding valuation of security investments.

In a similar study, Khansa and Liginlal [10] attempted to measure the level of investment in security by using the reported revenue and stock prices of the market leaders in the information

* Correspondence to: Department of Engineering Management and Systems Engineering, The George Washington University, Washington DC 20052, USA. Tel.: +1 202 994 7424; fax: +1 202 994 0245.
E-mail address: mazzu@gwu.edu (T.A. Mazzuchi).

0305-0548/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cor.2010.11.013

security product market as input variables while considering the monetized impact of highly publicized malicious software attacks (e.g., only network enabled virus, Trojan horse, and worm incidents). The basis of the monetized equivalencies was the data provided by the recognizably flawed CSI/FBI surveys of 1998–2006 (see Ref. [11]). While the authors claim that their analysis shows that they "demonstrated that with higher investment in information security comes more protection and resilience to malicious attacks", what they really showed was that the marketing efforts of the large security product firms have been successful. Beyond that, no meaningful conclusions are possible from this study.

The challenge in determining a value for information security measures lies in the fact that success is measured through events not occurring, or at least occurring infrequently, a situation which might have resulted from chance rather than effective security. Naturally, the attention of the information security profession is focused on trying to quantify probabilities of events occurring, which would provide managers with a way to compare their incident rates with a known probability of incidents. This would inform and enable the ability to use annualized loss expectancy (ALE) or other well-understood risk-impact equations as a measure of value in security investments, which is not possible without such information.

There have been a variety of attempts to gather a sense of what is occurring in security-related incidents. One, the Honeynet Project [12], relies on a set of computer systems that were specifically deployed to attract external attacks. The use of Honeynet data is extremely important as each participating computer in the Project is instrumented and used only to collect information on attacker methods over the internet. The data are extremely good for understanding network attack patterns and attacker techniques. It is not good; however, at providing insight into the parts of information security that it is not designed to address: those resulting from non-networked based attacks, such as insider abuse of access, theft of hardware, destruction of system components, or compromise of legitimate login credentials.

Other research efforts in the development of risk estimations have attempted to collect empirical data to be used in the characterization of risk in security. Karabacak and Sogukpinar [13] developed what they term the Information Security Risk Analysis Method (ISRAM), which is a survey-based tool to collect incident rates for use in standard risk calculations. The method is mathematically sound but limited by the need for incidents to have been both detected and understood to be incidents. As noted previously, neither of these characteristics is necessarily always true. Clever attackers are always attempting to find ways to attack without being detected and there is reason to believe that some, at least, have been successful. Certainly, vulnerability researchers such as Miller [14] specialize in finding what are called "zero day" vulnerabilities: vulnerabilities for which solutions are not available. These types of vulnerabilities have a great deal of value in the security community (on both sides).

Bodin et al. [15] used the Analytic Hierarchy Process (AHP) to weight order the elements of a composite metric for risk analysis. The composite metric consists of three elements: expected loss, expected severe loss, and standard deviation of loss. The use of this composite metric is intended to assist in assessing the risk associated with proposed solutions to security challenges. This assumes, again, that the ability to characterize the risk potential in terms of loss expectancy is possible from any reasonable perspective. As discussed above, this is not a given ability.

Additionally, there are several annual surveys conducted where respondents are asked to identify numbers, types, and impacts of incidents. Baker et al. [16] proposed an event-chain risk management model in which threats are "measured as rates per year and then converted into outcomes by specifying the number or extent per year." Again, this is limited because of the focus on threats that are both known and measurable, a dubious proposition at best. Sigonen et al. [17] focused less on counting threat incidents and more on measuring compliance with security policies as a surrogate variable for risk estimation. In some sense, this may be a more appropriate measure in that there is some basis in belief that compliance with security policies may reduce the exposure of an enterprise to security problems.

All of these efforts provide some insight, but none provides the comprehensive data required for a scientific treatment of the problem.

Gathering the probabilities using an empirical approach is extremely problematic, however. Simply collecting the number of incidents that are detected, through any detection means, by definition measures only a percentage – an unknown percentage – of the true incident rate. This is true because there is no way to know how many incidents were not detected, nor what their impacts were. Putting a captive system in harm's way and collecting information on the number of incidents that occur on that system only reflects the detectable attacks on the system. This provides only a little insight into the incident rate for highly dynamic interactive systems used by people for business processes.

The research reported in this article was an attempt at a novel approach to develop and to quantify a probability model associated with security-related incidents. In this research, the method of expert judgment was used to provide the quantification. Expert judgment has been successfully applied in many fields (see for example, [18]); however, this is the first time it has been used in the field of information security.

The research was carried out in the Washington, D.C., (USA) metropolitan area. Three groups participated: a control group, a set of experts, and a group of students studying information security. This paper deals with the analysis of the experts' contributions.

2. An overview of a model for cyber attacks

When trying to model a phenomenon whose probabilistic structure changes over time, a natural model to use is a stochastic process. Stochastic process models can range from very basic to very complex. A particularly robust and yet mathematically tractable class of stochastic processes is the Nonhomogeneous Poisson Process (NHPP). Readers interested in a full mathematical treatment are referred to Çinlar [19]; here only the basic properties will be presented. A NHPP is defined by a rate of occurrence of events, $m(t|\theta)$, which is often a parametric function depending on the set of parameters θ. For the problem at hand, these occurrences would be cyber attacks. Suggestions for such rates are given in Fig. 1 below—these include a constant rate of attack (referred to as a Homogeneous Poisson Process or HPP model), a decreasing rate of attack, an increasing rate of attack, and an attack rate that increases to a peak and then decreases. The attack rate modeling possibilities are limitless.

Given a rate of occurrence, the NHPP, $N(t)$ denoting the number of occurrences in the interval $[0,t]$, has the following useful properties

(P1) The mean or expected number of occurrences in $[0,t]$ denoted $M(t|\theta)$ is given by $M(t|\theta) = E[N(t)|m(t|\theta)] = \int_0^t m(u|\theta)\,du$

(P2) The probability mass function for the number of occurrences in $[0,t]$ is given by $Pr\{N(t) = n|M(t|\theta)\} = [(M(t|\theta))^n/n!]e^{-M(t|\theta)}$
Additionally for any time $s < t$

(P3) The Mean or Expected Number of Occurrences in $[s,t]$ is $M(t|\theta) - M(s|\theta)$

(P4) The Probability Mass Function for the Number of Occurrences in $[s,t]$ is given by $Pr\{N(t) - N(s) = n|M(t|\theta)\} = (([M(t|\theta) - M(s|\theta)]^n)/n!)e^{-[M(t|\theta)-M(s|\theta)]}$

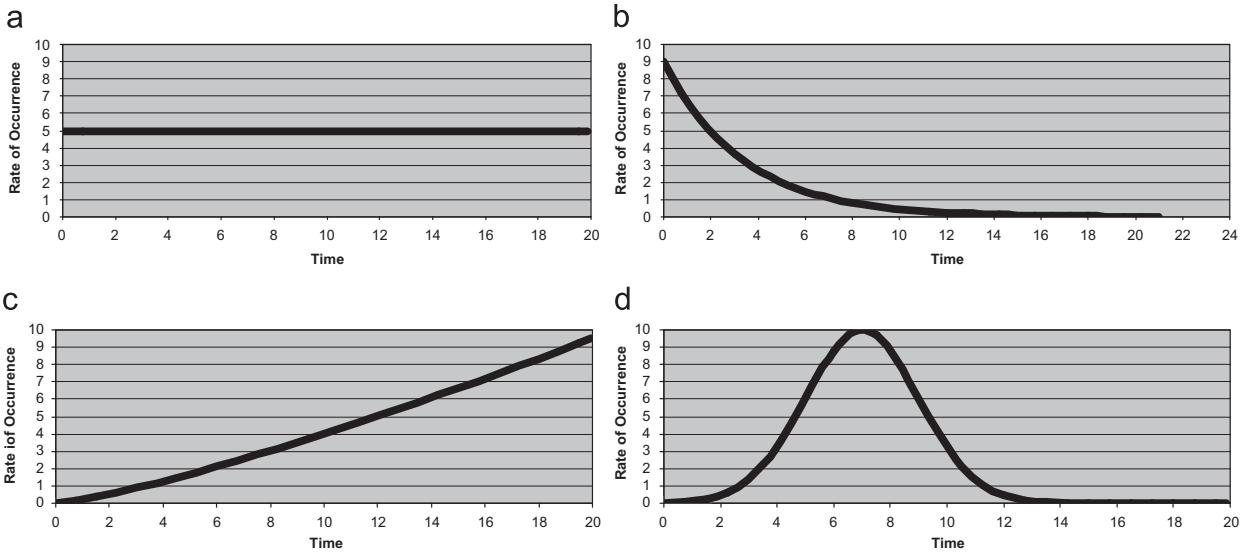

Fig. 1. Examples of rate of occurrence. (a) Constant: $m(t|\beta)=\beta$, (b) Exponential Decay: $m(t|\beta,\ \phi)=\beta\phi e^{-\phi t}$, (c) Power: $m(t|\beta,\ \phi)=\beta\phi t^{\phi-t}$, (d) Phase: $m(t|\beta,\mu,\phi)=(\beta/(\sqrt{2\pi}\phi))$ $e^{-(1/2)((t-\mu)/\phi)^2}$.

(P5) Given $M(t|\theta)$ and $s<t$, $N(s)$ and $N(t)-N(s)$ are independent random variables and thus $\Pr\{N(t)=n,N(s)=k|M(t|\theta)\}=\Pr\{N(s)=k|M(t|\theta)\}P\{N(t)-N(s)=n-k|M(t|\theta)\}$

(P6) Given two independent NHPPs $N_1(t)$ and $N_2(t)$ with respective mean value functions $M_1(t|\theta)$ and $M_2(t|\theta)$, the superimposed process given by $N_1(t)+N_2(t)$ is a NHPP with the mean value $M_1(t|\theta)+M_2(t|\theta)$

(P7) Given two independent NHPPs $N_1(t)$ and $N_2(t)$ with respective mean value functions $M_1(t|\theta)$ and $M_2(t|\theta)$,

$$Pr\{N_1(t)=k|N_1(t)+N_2(t)=n,M_1(t|\theta),M_2(t|\theta)\} = \binom{n}{k}\left(\frac{M_1(t|\theta)}{M_1(t|\theta)+M_2(t|\theta)}\right)^k\left(\frac{M_2(t|\theta)}{M_1(t|\theta)+M_2(t|\theta)}\right)^{n-k}$$

Often the functional form of $M(t|\theta)$ can be specified but the parameters θ must be determined through statistical inference procedures such as maximum likelihood or Bayesian estimation. Maximum likelihood estimation is performed by defining the likelihood function for observed data and estimating the value of the parameters θ that maximize the likelihood. The likelihood function for a NHPP process given the occurrence times $T_1, \ldots, T_n$, observed in an interval $[0,T^*]$, denoted $L(\theta|T_1, \ldots, T_n)$ is given by

$$L(\theta|T_1,\ldots,T_n)\propto \prod_{i=1}^{n} m(T_i|\theta)e^{-M(T^*|\theta)} \tag{1}$$

Thus, the maximum likelihood technique requires reliable observable data often in large quantities. This is generally not the case in a risk assessment situation and certainly not the case for the problem at hand.

Bayesian estimation is performed by defining a joint prior distribution for the parameters θ, $g(\theta)$, and estimating θ as the joint mean, median or mode of the distribution. This estimation procedure does not require data, but rather specification of the prior distribution through expert judgment. However, if data becomes available, a combined estimate is obtained by replacing the prior distribution by the posterior distribution, which is defined as

$$g(\theta|T_1,\ldots,T_n)\propto L(\theta|T_1,\ldots,T_n)g(\theta) \tag{2}$$

One advantage of Bayesian estimation is that predictive inference is straightforward, for example without observing any data

$$E[N(t)]=\int E[N(t)|M(t|\theta)]g(\theta)\,d\theta=\int M(t|\theta)g(\theta)\,d\theta \tag{3}$$

and

$$Pr\{N(t)=n\}=\int Pr\{N(t)=n|M(t|\theta)]g(\theta)\,d\theta \tag{4}$$

and given data, predictive inference is obtained as the above with the posterior distribution rather than the prior distribution. The disadvantage of Bayesian inference is that it requires the specification of a prior distribution. This often requires the elicitation, codification, and combination of expert judgment, which can often be a daunting set of tasks. A procedure for accomplishing such a task is discussed in the next section.

3. Expert judgment analysis

Expert judgment analysis is designed to elicit, codify, and combine the knowledge of people who have significant experience or expertize in a defined field in order to assess unknown quantities or parameters. It has been used in such wide ranging endeavors as assessing risks to the Washington State Ferry System [20], to nuclear power plant components [21], and to genetically modified crops [22]. The appropriate use of expert judgment is justified when quantitative data is missing, of dubious quality, or is insufficient for obtaining reasonable statistical results.

There are many methods for eliciting, codifying, and combining expert judgment (see for example, [18]). In combining expert judgment, often a weighted average of expert inputs is used and much research has centered on determining appropriate weighting schemes. For this research effort, the classical model of Cooke [18] was chosen. The reason for selection of this model is based on its extensive application in the field of risk analysis (see for example, [23]) and its demonstrated superior performance over the often used simple averaging technique (see for example, [24,25]). The model is summarized below, while a more thorough treatment is given in Ref. [18].

In the classical model, experts provide a distribution for unknown quantities by specifying 5th, 50th and 95th percentile values for the quantities of interest. The combination of the expert judgment is obtained as a convex combination of the expert distributions where the experts' weights are derived from the experts' responses to a set of seed variables whose values are known by the analyst and which are used to "calibrate" the accuracy of the experts' opinions. There are several advantages to this approach over the equal weighting of expert judgments.

3.1. The classical model

The process of expert judgment elicitation in the classical model follows several steps. First, the elicitation is prepared and experts are identified and selected. Then each expert is interviewed, alone and without knowledge of other experts to be interviewed or their responses. Experts are asked to specify 5th, 50th, and 95th percentile values for a series of quantities, some of direct interest and some related to the seed values used to calculate the weights for combining the expert judgment. Finally, the data are analyzed.

The analysis proceeds in two phases. First, an analysis is performed on the expert responses to the seed or calibration variables. A weight is assigned to each expert based on the expert's calibration and information scores. Calibration relates to how well the expert has specified his/her percentiles. A well calibrated expert will have approximately 5% of the seed variable realizations lower than his/her 5th percentile values, 45% of the seed variable realizations between his/her 5th percentile and 50th percentile values, 45% of the seed variable realizations between his/her 50th percentile and 95th percentile values, and 5% of the seed variable realizations greater than his/her 95th percentile values. Referring to these intervals as interval 1 through 4, the calibration score is calculated for the seed variables using *relative information, I*(*s*,*p*), given by

$$I(s,p) = \sum_{i=1}^{4} s_i \ln(s_i/p_i) \tag{5}$$

where s_i is the observed relative frequency of interval i from the seed variables and p_1, p_2, p_3, and p_4 are the values .05, .45, .45, and .05, respectively. It can be shown that when the number of seed variables, N, is reasonably large, the value $2NI(s,p)$ has a chi-squared distribution with 3 degrees of freedom, $\chi^2_{(3)}$. The calibration score, *cal*, for each expert is calculated as

$$cal = \begin{cases} Pr\{\chi^2_{(3)} > 2NI(s,p)\} & \text{if } Pr\{\chi^2_{(3)} > 2NI(s,p)\} > \alpha \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

That is, a minimum acceptable calibration score α is specified, otherwise the expert receives zero weight in the analysis.

In order to calculate the information score for each expert, the complete expert distribution must be defined. This is obtained by taking, for each seed variable, the known value, say r, and the elicited 5th, 50th, and 95th percentile values, say $q_{5,j}$, $q_{50,j}$, and $q_{95,j}$, respectively for each expert j, and defining 0th and 100th percentile point for the experts. This accomplished by defining

$$q_L = \min\{r, q_{5,1}, \ldots, q_{5,e}\}$$
$$q_U = \max\{r, q_{95,1}, \ldots, q_{95,e}\} \tag{7}$$

where e is the number of experts and then specifying a k% overshoot below q_L and above q_U for estimating the 0 percentile and 100 percentile point, respectively

$$q_0 = q_5 - (k/100)(q_U - q_L)$$
$$q_{100} = q_{95} + (k/100)(q_U - q_L). \tag{8}$$

The value of k is often set at 10. A complete expert distribution of the quantity is generated via linear interpolation between the five points as illustrated in Fig. 2. Calculated in this manner, the expert's distribution is minimally informative with respect to the uniform background measure. That is, we obtain a distribution which adheres to the expert's specified percentiles but is uniformly distributed between the percentiles thus imposing no additional assumptions. For the quantities of interest (that is, not including the seed variables), the bounds of the expert distributions are calculated using (7) and (8) excluding the value r (since these are the questions of interest as opposed to seed questions, there is no known realization).

The information score measures the deviation of the expert's distribution with respect to some meaningful background measure, which in this case is taken to be the uniform distribution over the entire range $[q_0, q_{100}]$. (see Fig. 3). As such, it is a measure of the expert's certainty in his/her answers (see for example, [26]). The information measure is calculated again using relative information

$$\text{inf} = I(h,p) = \sum_{i=1}^{4} p_i \ln(p_i/h_i) \tag{9}$$

where

$$h_1 = F_U(q_5) - F_U(q_0) = \frac{q_5 - q_0}{q_{100} - q_0}$$
$$h_2 = F_U(q_{50}) - F_U(q_5) = \frac{q_{50} - q_5}{q_{100} - q_0}$$
$$h_3 = F_U(q_{95}) - F_U(q_{50}) = \frac{q_{95} - q_{50}}{q_{100} - q_0}$$
$$h_4 = F_U(q_{100}) - F_U(q_{95}) = \frac{q_{100} - q_{95}}{q_{100} - q_0} \tag{10}$$

and F_U is the CDF of the uniform background measure. The information score for each expert can be calculated as the average seed variable information score or can be calculated individually for each quantity of interest. An overview of the concepts of these scores is presented in Fig. 4.

The final weight for each expert is determined as the normalized (weights sum to one) product of the calibration score and the information score. The determination of the value of α is selected so that a fictitious expert, whose distribution would be the resulting

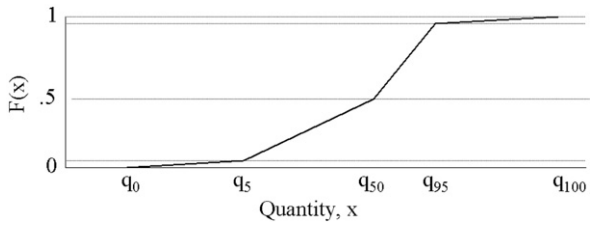

Fig. 2. Example expert distribution for unknown quantity.

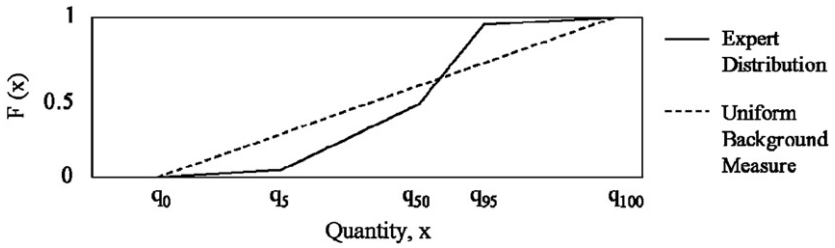

Fig. 3. Comparison of expert distribution with background measure.

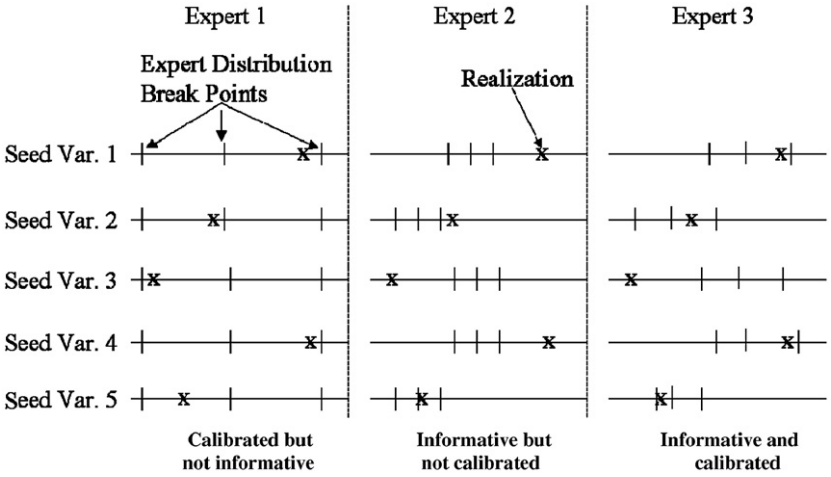

Fig. 4. An Illustration of calibration and information concepts.

weighted mixture of the expert distributions, would receive the highest possible weight if added to the expert pool.

The second phase of the analysis addresses the actual quantities or parameters of interest. In this phase, the experts' distributions for the quantity of interest are combined using weights obtained.

4. The research

The purpose of this research was to use expert judgment elicitation to provide data to develop a model for evaluating the following questions:

- How often does a computer or system come under attack?
- How many of those attacks are successful?
- It is worthwhile to make a large investment to protect a system from attacks?
- How probable is a successful attack under different protection scenarios?

These questions are problematic in the information security realm because of several complications. First, success is defined as nothing bad happening, or at least bad things happening relatively infrequently. But if no (few) attacks are detected, does that really mean nothing bad is happening? Or could it mean that attacks are occurring but are not being detected? Or could it simply be coincidence that no attacks are observed? Second, the challenge of measuring the number of attacks is daunting. In a single enterprise, the best that can be determined is how many attacks have been detected; however, it is difficult to know what percentage of the total number of attacks this represents. Within a large interconnected information infrastructure consisting of many enterprises, the data may not be available to determine even the total number of detected attacks. There are many strong imperatives for enterprises to not reveal detected (successful or otherwise) attacks. One of these is that it may encourage other attackers to choose that enterprise as a target. Another is that revealing such problems could have a deleterious effect on the perceptions of the customer base for the enterprise which could result in criminal or civil legal or regulatory liability, damaged reputation, loss of customer confidence, reduction in stock prices, or loss of market share.

There are many examples. A serious data breach at the retailer TJ Maxx began in 2005 and continued until 2007. One or more hackers stole 45.7 million credit card and debit card numbers that had been used by customers for purchases. The hackers also stole Social Security numbers, driver's license numbers, and military identification numbers. Fraud based on the stolen information occurred in at least seven States and eight foreign countries. The total impact on TJ Maxx is unknown but estimates are that an amount exceeding US$1 billion will be needed for security upgrades, consultants, legal fees, and public relations, but not including legal liabilities. Upwards of 21 lawsuits have been filed seeking damages from the company that will increase the loss (see for example, [27]).

In another example, Hewlett-Packard's Chairwoman faced criminal charges following her response to a leak of information by a director of the company. She hired a team of security professionals to spy on the communications in the personal accounts of ten other directors, including both email and telephone calls from their homes and private cell phones. She found the source of the leak, but he refused to resign. Another director, however, did resign, triggering a requirement for an 8-k report to the SEC, but the initial filing ignored the SEC requirement to report a disagreement regarding operation, policies or practices of the company. The impacts extend beyond the boardroom and Patricia C. Dunn's suitability for the Chair—she was ousted in 2006. Civil lawsuits for invasion of privacy, intentional infliction of emotional distress, and unfair business practices have been filed (see for example, [28–30]) Dunn and four others have been charged with criminal fraud and conspiracy (see for example, [31]).

In the information security business, 30 years of practical and theoretical research has resulted in a fairly sophisticated appreciation for how to judge the qualitative level of risk faced by an enterprise. This includes how to identify and mitigate vulnerabilities, how to identify and constrain threats, how to identify appropriate countermeasures, and sometimes how to quantify the exposure factor related to the impact of a successful attack on a specified information asset or system. Based upon that understanding, there is a practical level of protection that a competent security manager can architect for a given enterprise. In other words, if the responsible manager is appropriately educated and adequately funded, it is possible to take reasonable steps to protect information assets and systems, even if the information that would permit optimization is not available.

However, not all managers are equally educated or experienced, nor are all adequately funded. There remains a continual challenge to explain and defend security expenditures to higher-level managers, who generally demand a business case analysis. Making this business case based on qualitative data is difficult. Unfortunately, quantitative data does not exist. There have been many attempts to develop quantitative data using traditional quantitative methods, such as experiments, surveys, and observations, but

there are significant weaknesses apparent in each approach. With experiments, the resultant data to date only reflects the controlled responses to known and scripted attacks. With surveys, the data only reflects the knowledge and awareness of the respondents. Furthermore, many of the surveys have significant structural weaknesses that subvert any derived meaning or utility for quantitative analyzes (see for example, [11]). Quantitatively derived models only approximate what reality might truly be, and do that reflecting the biases or limitations of the model developers.

This research was constructed to explore the utility of applying the well-established method of expert judgment elicitation to the field of information security. The instrument for eliciting the expert judgments was developed by two information security specialists and two expert judgment analysis specialists. The resultant instrument was validated using a small set of information security experts. The final instrument was used to elicit answers to both the calibration (seed variable) and judgment questions through structured interviews. An expert judgment analysis specialist conducted the interviews. The data were compiled and analyzed by the specialist in expert judgment analysis. In all, the research was conducted over the period of six months.

5. The research instrument

Experts were asked a series of 31 questions broken up into four sections. The first section contained 10 seed variable questions. The questions, presented in Table 1 with their realizations, were developed by the two information security specialists with the notion that information security experts did not have direct access to the actual answer but had adequate knowledge to make the required assessments.

Sections 2 through 4 contain the same six questions but the expert was asked to consider the questions under the 3 separate scenarios presented in Table 2 below. These scenarios represent the best possible, S_{BP}, the most likely, S_{ML}, and the "honey pot", S_{HP}, scenarios. Sections 2 and 3 also contained a question about the obsolescence time of security software. These questions, presented in Table 3, form the basis of the analysis.

For questions Q_1, Q_2, Q_3, Q_4 and Q_7, experts were permitted to use any time units that they felt comfortable with, however, all values were converted to hours in the analysis.

Table 1
Seed variable questions and their realizations.

Question	Realization
If today you are using a 128-bit cryptographic key, how many bits would you need to use in six years in order to maintain the same level of security as you enjoy today?	130
How likely is an attack on a system to be successful?	0.65
If a system is successfully attacked, how likely is it that the attack will be detected?	0.04
How likely is a successful attack on a system to be detected and reported to authorities by managers responsible for the security of the system?	0.01
How long would it take a brute force attack by a single computer on a 56-bit cryptographic key to recover the key? (Hours)	72
How long would it take a brute force attack by a distributed network of computers on a 56-bit cryptographic key to recover the key? (Hours)	22
How much more attack activity can we expect this year than last in attack per company per week?	0.2
How much more likely, if at all, is a company with more than 500 employees to be attacked than a company with less than 500 employees?	0.5
How much more often are public companies attacked than private or not-for profit companies?	2
What percentage of vulnerabilities in computer systems and networks are easy to exploit, requiring only moderate computer skills?	70%

Table 2
Section II–IV scenario descriptions.

Section	Notation	Scenario description
II	S_{BP}	This section addresses your expectations given that all reasonable steps have been taken to protect information assets and systems from attack—that is, that the protection is comparable to the best available protection for systems and networks that is in use today and competently managed.
III	S_{ML}	This section addresses your expectations given that average or usual care has been taken to protect information assets and systems from attack—that is, that the protection is comparable to what you would expect to find in most systems and networks in use today.
IV	S_{HP}	This section addresses your expectations given that information assets and systems are competently managed but do not employ any protection systems.

Table 3
Section II–IV scenario questions.

Notation	Scenario II, III, and IV questions
Q_1	How long do I have to wait before the first attack?
Q_2	How long do I have to wait after the first attack before the second attack occurs?
Q_3	How long do I have to wait before the first successful attack?
Q_4	How long do I wait after the first successful attack before the second successful attack occurs?
Q_5	How many attempted attacks can you expect in one month?
Q_6	How many successful attacks can I expect in one month?
	Additional Question for Scenario II and III
Q_7	How fast does the security software become obsolete?

6. Expert judgment analysis

The analysis of expert judgment using the classical model was performed using Microsoft Excel . Fig. 5 displays "range graphs" for the 13 experts' responses to several seed variables. The 5th and 95th percentile values supplied by the experts are denoted by "|" and their median values by "X". The calculated upper and lower bounds are denoted by "•" and the realization by a dashed line. Note that the bounds calculated in Eq. (8) were truncated to represent reality. That is, for example, questions regarding probability values were restricted to [0,1]. The potential problem with equal weights can be seen in Fig. 5 where small groups of experts have specified a significantly different range of uncertainty for the seed variables. An equally weighted combination of expert distributions could thus tend to have more spread or be less informative. In the classical method, experts who habitually provide less informative assessments are typically given less weight.

Based on the expert responses to the seed variables, the calibration and information scores were obtained for the experts and are presented in Table 4 below with weights determined by these scores. Experts were given IDs to preserve anonymity. The range of calibration is given by expert 8 with the lowest score to experts 11 and 12 with identical high scores. The range of informativeness is given by expert 4 with the lowest score to expert 5 with the highest score. It can be seen that the range of calibration scores is more pronounced that that of the information

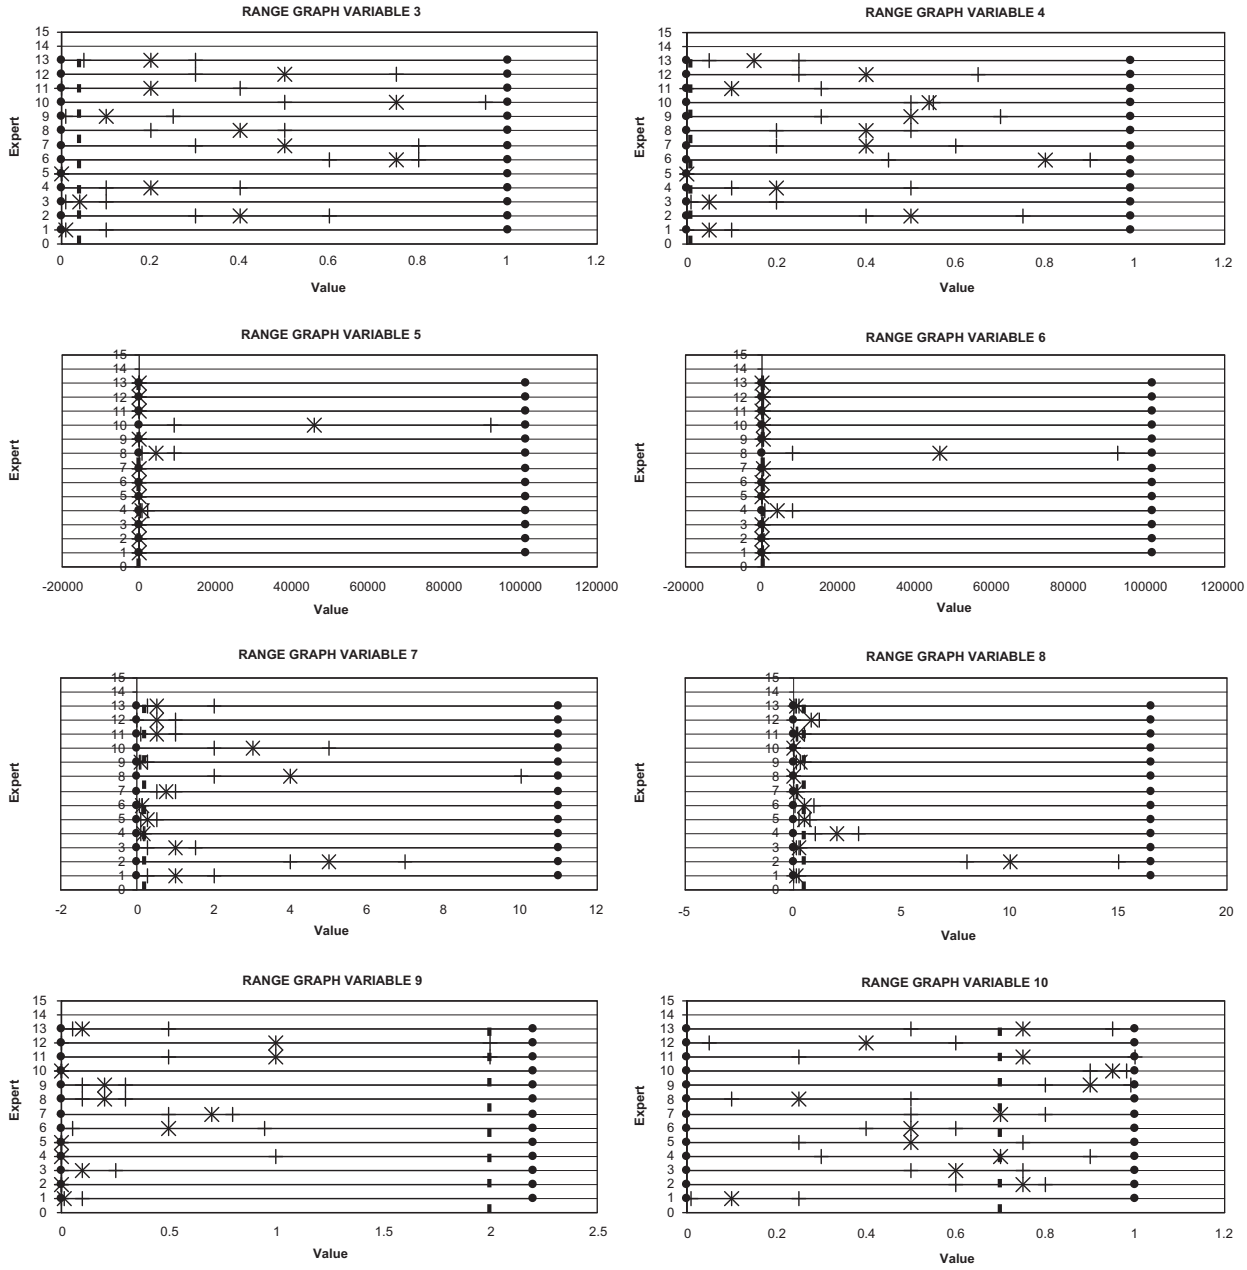

Fig. 5. Range graphs for seed variable questions.

Table 4
Expert scores for seed variables.

Expert ID	Calibration score	Information score	Unnormalized weight	Normalized weight
1	4.559E−04	3.4271	0.000	0.000
2	1.362E−06	4.0296	0.000	0.000
3	1.397E−02	2.7781	0.000	0.000
4	4.559E−04	1.9192	0.000	0.000
5	2.083E−05	7.1414	0.000	0.000
6	1.066E−06	3.0101	0.000	0.000
7	1.102E−03	2.7624	0.000	0.000
8	5.448E−08	2.0835	0.000	0.000
9	4.704E−02	2.4445	0.000	0.000
10	4.078E−06	3.6292	0.000	0.000
11	3.135E−01	2.4765	0.776	0.543
12	3.135E−01	2.0829	0.653	0.457
13	1.579E−05	2.6472	0.000	0.000

score. Generally speaking, calibration is more important than informativeness. The experts who stood out as having both good informativeness scores and high calibration values were experts 11 and 12. All other experts received weight 0. It is not unusual for this technique to assign a zero weight to a sizeable number of experts.

The expert and combined expert percentile values for the questions of interest are presented in Table 5. Several observations are worth noting. First, the expert distributions for questions Q1, Q2, Q5, Q7 are identical regardless of the scenario. This would be expected as it refers to external behavior not affected by the change in the scenario. This result shows consistency among the experts. Expert distributions for questions Q3, Q4, and Q6 do differ and display the proper ordering as we move from the more protected to a less protected scenario.

Another observation worth noting is that the distance between the scenario 1 distributions and distributions for scenarios 2 and

Table 5
Expert and combined expert percentiles and the questions of Table 3.

Question	How long do I have to wait before the first attack (in hours)?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	1.60E−02	8.30E−02	1.00E+00	1.60E−02	8.30E−02	1.00E+00	1.60E−02	8.30E−02	1.00E+00
12	5.55E−04	2.40E+01	7.20E+01	5.55E−04	2.40E+01	7.20E+01	5.55E−04	2.40E+01	4.80E+01
Combined	1.59E−02	8.28E−01	7.20E+01	1.59E−02	8.28E−01	7.20E+01	1.59E−02	8.28E−01	4.80E+01
Question	How long do I have to wait after the first attack before the second attack occurs (in hours)?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	1.60E−02	8.30E−02	1.00E+00	1.60E−02	8.30E−02	1.00E+00	1.60E−02	8.30E−02	1.00E+00
12	1.38E−03	3.00E+01	1.20E+02	1.38E−03	3.00E+01	1.20E+02	1.38E−03	3.00E+01	1.20E+02
Combined	1.59E−02	8.33E−01	1.20E+02	1.59E−02	8.33E−01	1.20E+02	1.59E−02	8.33E−01	1.20E+02
Question	How long do I have to wait before the first successful attack (in hours)?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	2.40E+01	7.68E+02	9.22E+03	2.40E+01	1.92E+02	3.84E+02	1.66E−01	3.33E−01	5.00E−01
12	4.80E+01	3.84E+02	7.68E+02	1.00E+00	8.00E+00	7.20E+01	4.80E+01	1.92E+02	5.04E+02
Combined	3.82E+01	5.30E+02	9.15E+03	1.88E+00	6.54E+01	3.84E+02	1.82E−01	4.89E−01	4.96E+02
Question	Hong long do I have to wait after the first successful attack before the second successful attack occurs (in hours)?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	2.40E+01	7.68E+02	9.22E+03	2.40E+01	1.92E+02	3.84E+02	1.66E−01	3.33E−01	5.00E−01
12	7.20E+01	7.68E+02	1.54E+03	2.40E+01	1.20E+02	2.40E+02	7.20E+01	3.36E+02	7.68E+02
Combined	4.76E+01	7.68E+02	9.16E+03	2.40E+01	1.53E+02	3.84E+02	1.82E−01	4.89E−01	7.59E+02
Question	How many attempted attacks can you expect in one month?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	1.00E+03	1.00E+04	2.00E+04	1.00E+03	1.00E+04	2.00E+04	1.00E+03	1.00E+04	2.00E+04
12	5.00E+00	3.50E+01	1.00E+02	5.00E+00	3.50E+01	1.00E+02	5.00E+00	3.50E+01	1.00E+02
Combined	8.93E+00	2.43E+03	2.00E+04	8.93E+00	2.43E+03	2.00E+04	8.93E+00	2.43E+03	2.00E+04
Question	How many successful attacks can I expect in one month?								
Scenario	1			2			3		
Expert	5%	50%	95%	5%	50%	95%	5%	50%	95%
11	0.00E+00	1.00E+00	5.00E+00	5.00E+00	1.00E+01	2.00E+01	5.00E+02	1.00E+03	2.00E+03
12	0.00E+00	2.00E+00	7.00E+00	0.00E+00	5.00E+00	1.00E+01	1.00E+00	3.00E+00	1.00E+01
Combined	0.00E+00	1.63E+00	7.00E+00	5.83E−01	7.72E+00	2.00E+01	1.26E+00	5.79E+02	2.00E+03
Question	How fast does the security software become obsolete (in hours)?								
Scenario	1			2					
Expert	5%	50%	95%	5%	50%	95%			
11	4.61E+03	9.22E+03	2.77E+04	4.61E+03	9.22E+03	2.77E+04			
12	4.61E+03	9.22E+03	1.84E+04	4.61E+03	9.22E+03	1.84E+04			
Combined	4.61E+03	9.22E+03	2.65E+04	4.61E+03	9.22E+03	2.65E+04			

3 is more pronounced than the distance between the distributions for scenario 2 and 3. A first order conclusion is that there are benefits from even some protection. The marginal return for increasing levels of protection is not as well pronounced. This analysis should be the subject of future research.

7. Using expert judgment results

The stated purpose of this research was to illustrate the development of prior distributions for the parameters of NHPP models for cyber attacks, thus providing a fully quantified model for use in risk analysis. The expert judgment results from the previous section can be used to develop just such distributions for both HPP and NHPP forms.

7.1. HPP for successful and unsuccessful attacks

For the first model, it will be assumed that the rate of successful and unsuccessful attacks (denoted β_{Success} and $\beta_{\text{Unsuccess}}$, respectively) remains constant over time as in model "a" in Fig. 1. That is, the number of successful and unsuccessful attacks, respectively, are given by an HPP and thus via (P6) the number of attacks is also

an HPP with rate β_{Attack}. For this model, prior distributions for parameters for the attack and successful attack processes are directly obtained from Q5 and Q6. That is, Q5 addresses the average number of occurrences per month for the combined successful and unsuccessful attacks and Q6 addresses the average number of occurrences of successful attacks per month. Prior best estimates for β_{Attack} and β_{Success} for each scenario can be directly obtained as the median estimate for question Q5 and Q6, respectively, and are presented in Table 6.

Several additional quantities may be obtained through simulation. For example, the probability distribution for the probability of a successful attack may be obtained by simulation using the distributions for β_{Attack} and β_{Success} for each scenario and calculating, using (P7)

$$Pr\{\text{Successful Attack}\} = \frac{\beta_{\text{Success}}}{\beta_{\text{Attack}}} \tag{11}$$

Simulation using the expert distributions is quite simple as it has finite support and the cumulative distribution is a linear interpolation between points. Noting that $\beta_{\text{Attack}} \geq \beta_{\text{Success}}$ when simulating (11) above, the procedure would be to first simulate β_{Attack} and then simulate β_{Success} from the conditional distribution

$$Pr\{\beta_{\text{Success}} < \beta | \beta_{\text{Success}} < \beta_{\text{Attack}}\} \tag{12}$$

An example calculation is given in Table 7 below.

Probability distributions for the number of attacks or the number of successful attacks within a specified time frame may be similarly obtained from (P4) through simulation by using the appropriate distribution of β and calculating

$$Pr\{N(t)-N(s) = n\} = \frac{(\beta(t-s))^n}{n!} e^{-\beta(t-s)} \tag{13}$$

for the probability that the number of attacks in the interval $[s,t]$ is n. Example calculations for the probability distribution for the probability of no successful attacks in a month is provided in Table 8.

Table 6
Estimates of β for successful and unsuccessful attack processes.

Attack	Scenario 1 success	Scenario 2 success	Scenario 3 success
2430.41	1.63	7.72	579.24

Table 7
Percentiles of the distribution of probability of successful attack.

Scenario	5%	50%	95%
1	0.00E−00	4.60E−04	1.14E−01
2	1.02E−05	2.00E−03	2.93E−01
3	4.91E−05	1.08E−01	8.21E−01

Table 8
Percentiles of the distribution of probability of no successful attacks in a month.

Scenario	5%	50%	95%
1	9.51E−04	1.98E−01	1.00E+00
2	2.17E−09	4.81E−04	5.64E−01
3	0.00E−00	1.09E−251	1.29E−01

7.2. NHPP for successful and unsuccessful attacks

Next we address the model where attack rates are time dependent. Although there are several available models, model "c" in Fig. 1 is selected for illustrative purposes. Assuming that both the total number of attacks and the number of successful attacks can be modeled with a power intensity NHPP where the parameters to be estimated are denoted (β_{Attack}, ϕ_{Attack}) and (β_{Success}, and ϕ_{Success}) for the number of attacks and number of successful attacks, respectively. However, the estimation is more difficult.

To construct the simulation, we obtain the best estimates from the experts of the times $T_1^{(A)}$ and $T_2^{(A)}$, the time until the first attack and the time between the first and second attack and of the times $T_1^{(S)}$ and $T_2^{(S)}$, the time until the first successful attack and the time between the first successful and second successful attack. These are obtained from the answers to Q1–Q4, respectively. These are equated to their theoretical mean values obtained from the NHPP with power intensity.

Consider a NHPP with rate $m(t)$. If T_i denotes the ith interarrival time and S_i denotes the ith arrival time, then the marginal distribution of S_n is given by van Noortwijk et al. [32] as

$$f_{S_n}(y) = \frac{1}{(n-1)!}\left[\int_0^y m(t)\,dt\right]^{n-1} m(y) e^{-\int_0^y m(t)\,dt} \tag{14}$$

Since it is assumed that $m(t) = \beta\phi t^{\phi-1}$ the respective densities of S_1 and S_2 may be calculated

$$\begin{aligned} f_{S_1}(y) &= \phi\beta y^{\phi-1} e^{-\beta y^{\phi}} \\ f_{S_2}(y) &= (\beta y^{\phi})\phi\beta y^{\phi-1} e^{-\beta y^{\phi}} \end{aligned} \tag{15}$$

The first distribution is the well known Weibull distribution and the expected value expression is

$$E[Y] = \beta^{(-1/\phi)}\Gamma\left(1+\frac{1}{\phi}\right) \tag{16}$$

Likewise it can be shown that the expectation for the second distribution can be expressed as

$$E[Y] = \beta^{(-1/\phi)}\Gamma\left(2+\frac{1}{\phi}\right) \tag{17}$$

Thus to simulate from the joint distribution of ϕ and β we use the following procedure

1. Simulate t_1 and t_2 from the experts' distribution for time until first and time between the first and second attack.
2. Calculate $s_1 = t_1$ and $s_2 = t_1 + t_2$.
3. Equate
$$\begin{aligned} s_1 &= E[S_1] = \beta^{(-1/\phi)}\Gamma(1+1/\phi) \\ s_2 &= E[S_2] = \beta^{(-1/\phi)}\Gamma(2+1/\phi) = (1+1/\phi)s_1 \end{aligned}$$
4. Calculate
$$\phi = \frac{1}{((s_2/s_1)-1)}$$
$$\beta = \left(\frac{\Gamma(1+(1/\phi))}{s_1}\right)^{\phi}$$

The distribution of ϕ is important. If the distribution is concentrated about 1, that would indicate that a HPP was appropriate.

Simulated distributions for both the shape parameter, ϕ, and scale parameter, β, were obtained. Selected percentiles are given in Table 9. Both distributions show considerable variability which is desirable for a prior distribution in that it does specify overconfidence in the prior

Table 9
Percentiles of shape and scale parameters of the power intensity model.

Percentile	ϕ	β
5th	2.39E−04	0.00E+00
25th	2.26E−02	1.06E−22
50th	8.50E−01	4.60E+00
75th	3.64E+01	7.85E+01
95th	2.64E+03	7.36E+04

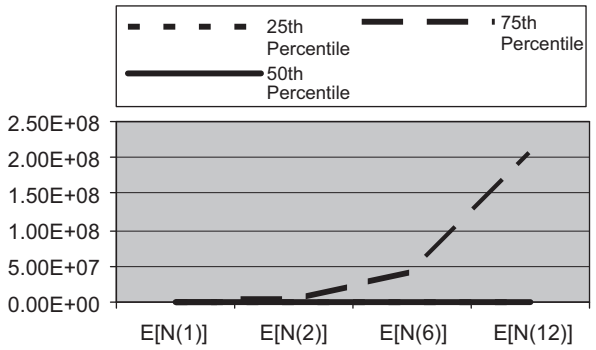

Fig. 6. Percentile plot for number of attacks as a function of time.

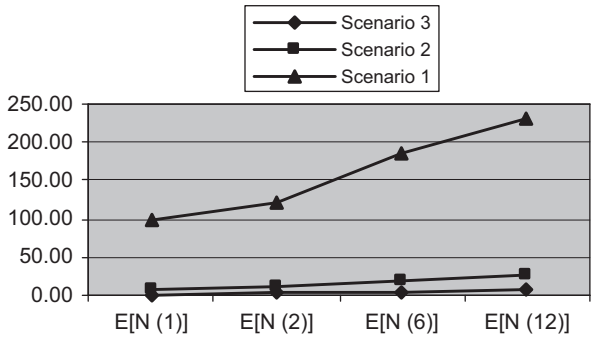

Fig. 7. Median value plot for number of successful attacks as a function of time for each scenario.

estimates. Most interesting is the distribution of the shape parameter which is not concentrated about 1. This would indicate that the NHPP is more appropriate than the HPP as a model.

Next we use the distribution to simulate the expected number of attacks as a function of time. Specifically we consider, for each scenario, the expected number of attacks and the expected number of successful attacks in 1, 2, 6, and 12 months, respectively, where 1 month = 28 days. Fig. 6 illustrates the usual increase in uncertainty as we predict further ahead in time. This figure is for the number of attacks but similar figures can be presented for the number of successful attacks under each scenario. Note that the 25th percentile is masked almost completely by the 50th percentile.

Fig. 7 compares the median number of successful attacks of scenarios 1, 2, and 3 as a function of time. As with the HPP results, it is apparent that at least some protection can make a considerable difference in the number of successful attacks. For example, it can be seen from Fig. 7 that the ratio of the number of successful attacks between scenario 1 and 2 is approximately 3. The ratio of the number of successful attacks between scenario 2 and 3 is an order of magnitude.

8. Conclusions

In information infrastructures that enjoy a sound, well-managed information security program, incidents are rare, and, consequently, statistical information that can be used for quantitative risk management evolves so slowly that it cannot keep up with the evolution of the threat environment. The research described in this paper was constructed to explore the utility of applying the well-established method of expert judgment elicitation to the field of information security, providing an alternative methodology for determining probabilities of successful attacks. The instrument for eliciting the expert judgments was developed by two information security specialists and two expert judgment analysis specialists. The resultant instrument was validated using a small set of information security experts.

Three scenarios were analyzed. The first was the case where all reasonable steps have been taken to protect information assets and systems from attacks (competent and protected case). The second was the case where average or usual care has been taken to protect information assets and systems (normal case). The third was the case where the information systems and assets have been competently managed but no special protection mechanisms are employed (competent but not protected). It is clear that the probability of a successful attack is smallest in scenario 1 but that the rate of attacks is higher than in the other scenarios. This matches well with our expectations for this scenario, since large enterprises like banks, governments, and military offices have very high rates of attacks but are generally well protected. The rate of attacks in scenario 2 is smaller than in the other two scenarios, but the rate of successful attacks is larger. For the third scenario, the lack of protection mechanisms may attract attacks but the well-managed aspect decreases the success rate somewhat.

This research illustrates the development of prior distributions for the parameters of models for cyber attacks and uses expert judgment results to develop the distributions.

In the very near term, a competently managed system with no protections is slightly less likely to observe a successful attack, but that advantage drops off over the longer term. The competently managed and well-protected system is always less likely to experience a successful attack as compared to the other two scenarios. The generalized conclusion that derives from this analysis is that when the attack rate is constant, an investment in information security protections will provide a decrease of approximately 99% in the average number of successful attacks when protection moves from competent management (scenario 3) to competent security (scenario 2). A further reduction of approximately 79% can be obtained by additional investments implementing excellent security (scenario 1). When the attack rate is modeled as power intensity NHPP, these percentages are 96% over a years time when protection moves from competent management (scenario 3) to competent security (scenario 2). Under this model, a further reduction of approximately 68% over the year's time can be obtained by additional investments implementing excellent security (scenario 1).

The decision on whether to invest or not to invest in information security protections is of course dependent upon the unique situation of an enterprise. However, for those enterprises with high exposure or valuable assets, this analysis confirms our intuition that reasoned investments in both competent security management and protection technologies would be rational decisions.

References

[1] Wright PD, Liberatore MJ, Nydick RL. A survey of operations research models and applications in homeland security. Interfaces 2006;36(6):514–29.
[2] Shindo S, Yamazaki H, Toki A, Mseshima R, Koshijima I, Umeda T. Approach to potential risk analysis of networked chemical plants. Computers and Chemical Engineering 2000;24(No. 2):721–7.

[3] Winkler RL, Wallsten TS, Whitfeld RG, Richmond HM, Hayes SR, Rosenbaum AS. An assessment of risk of chronic lung injury attributable to long-term ozone exposure. Operations Research 1995;43(No. 1):19–33.
[4] Erkut E, Verter V. Modeling of transport risk for hazardous materials. Operations Research 1998;46(5):625–42.
[5] Barnett A. Free flight and en route air safety: a first order analysis. Operations Research 2000;48(6):833–45.
[6] Dillon RL, Pate-Cornell ME, Guikema SD. Programmatic risk analysis for critical engineering systems under tight resource constraints. Operations Research 2003;51(3):354–70.
[7] Wu DD, Olson DL. Enterprise risk management: coping with model risk in a large bank. Journal of the Operational Research Society 2010;61(2):179–90.
[8] Wu DD, Kefan X, Gang C, Ping G. A risk analysis model in concurrent engineering product development, to appear. Risk Analysis 2010.
[9] Davis G, Garcia A, Zhang W. Empirical analysis of the effects of cyber security incidents. Risk Analysis 2009;29(9):1304–16.
[10] Khansa L, Liginlal D. Quantifying the benefits of investing in information security. Communications of the ACM 2009;52(11):113–7.
[11] Ryan JJCH, Jefferson TI. The use, misuse and abuse of statistics in information security research. In: Proceedings of the ASEM national conference, American society of engineering management, 2003; p. 644–653.
[12] The honeynet project. Know your enemy: learning about security threats, 2nd ed. Addison-Wesley Professional: Boston, 2004. See also 〈http://www.honeynet.org〉.
[13] Karabacak B, Sogukpinar I. ISRAM: Information security risk analysis method. Computers and Security 2005;24:147–59.
[14] Miller C. Kim Jong-il, Me. How to build a cyber army to attack the US. In: Proceedings of the conference on cyber conflict, NATO Center for Cooperative Cyber Defense Center of Excellence, Tallinn, Estonia, June 15–18, 2010.
[15] Bodin LD, Gordon LA, Loeb MP. Information security and risk management. Communications of the ACM 2008;51(No. 4):64–8.
[16] Baker WH, Rees LP, Tippett PS. Necessary measures: metric-driven information security risk assessment and decision making. Communications of the ACM 2007;50(10):101–6.
[17] Siponen MM, Mahmood A, Pahnila S. Are employees putting your company at risk by not following information security policies? Communications of the ACM 2009;52(12):145–7.
[18] Cooke RM. Experts in Uncertainty. Oxford University Press: Oxford, UK; 1991.
[19] Çinlar E. Introduction to Stochastic Processes. Prentice Hall: Upper Saddle River, NJ; 1975.
[20] Van Dorp JR, Merrick J, Mazzuchi TA, Harrald J, Grabowski M. A risk management procedure for the Washington state ferries. Risk Analysis 2001;21:127–42.
[21] Vo TV, Simonen FA, Gore BF, Livingston JV. Expert judgment elicitation on component rupture probabilities for five PWR systems. American Society of Mechanical Engineers, Pressure Vessels and Piping Division (Publication) PVP, v 251, Reliability and Risk in Pressure Vessels and Piping, 1993: p. 127–140.
[22] Krayer Von Krauss MP, Casman EA, Small MJ. Elicitation of expert judgments of uncertainty in the risk assessment of herbicide-tolerant oilseed crops. Risk Analysis 2004;24(6):1515–28.
[23] Cooke RM, Goossens LLHJ. TU Delft expert judgment data base. Reliability Engineering and System Safety 2008;93:657–74.
[24] Cooke RM, ElSaadany S, an Huang X. On the performance of social network and likelihood-based expert weighting schemes. Reliability Engineering and System Safety 2008;93:745–56.
[25] Cooke RM. Response to discussants. Reliability Engineering and System Safety 2008;93:775–7.
[26] Cooke RM, Goosesens LJH. Procedures guide for structured expert judgment, Nuclear Science and Technology. Delft University of Technology: Delft, the Netherlands, 2000.
[27] Pereira J. Breaking the code: how credit-card data went out the wireless door. The Wall Street Journal Online, Dow Jones & Company, Inc., 2007. See also http://online.wsj.com/article/SB117824446226991797.html〉.
[28] Kaplan D. Intrigue in high places. The Washington Post Company: Newsweek Business, 2006. See also 〈http://www.msnbc.msn.com/id/14687677/site/newsweek/〉.
[29] Darlin D. Journalists intend to sue Hewlett-Packard over surveillance. The New York Times, 2007. See also 〈http://www.nytimes.com/2007/05/07/business/media/07hp.html?ex=1189828800&en=47575ea1769e950d&ei=5070〉.
[30] Associated Press. HP sued in boardroom surveillance scheme. CBS Corporation: CBS News, 2007. See also 〈http://www.cbsnews.com/stories/2007/08/15/business/main3172338.shtml〉.
[31] Nagashima E, Noguchi Y. Dunn, four others charged in Hewlett-surveillance case. Washington Post.com, 2006. See also 〈http://www.washingtonpost.com/wp-dyn/content/article/2006/10/04/AR2006100401072.html〉.
[32] van Noortwijk J, van der Weide H, Kallen MJ. Technical report, Delft University of Technology, 2006.