

Contagion in Cybersecurity Attacks

Adrian Baldwin^a, Iffat Gheyas^b, Christos Ioannidis^c, David Pym^d, Julian Williams^{e,*}

^aResearcher, HP Labs, Bristol, Stoke Gifford, Bristol BS34 8QZ.

Email: adrian.baldwin@hp.com. Tel. +44 (0)117 3162400

^bLecturer in Big Data Analytics, Birmingham City University,

Email: iffat.gheyas@bcu.ac.uk. Tel. +44(0)191 334 2000

^cHead of Economics and Professor of Economics and Finance, Department of Economics,

University of Bath, BA2 3AY.

Email: c.ioannidis@bath.ac.uk. Tel. +44(0)1225 383226

^dProfessor of Information, Logic, and Security and Head of Programming Principles, Logic, and Verification, University

College London, Department of Computer Science, London WC1E 6BT

d.pym@ucl.ac.uk

^eProfessor of Finance, Durham University Business School, University of Durham, Durham DH1 3LE

Email: julian.williams@durham.ac.uk. Tel. +44(0)191 334 5301

Abstract

Systems security is essential for the efficient operation of all organizations. Indeed, most large firms employ a designated ‘Chief Information Security Officer’ (CISO) to coordinate the operational aspects of the organization’s information security. Part of this role is in planning investment responses to information security threats against the firms corporate network infrastructure. To this end, we develop and estimate a vector equation system of threats to ten important IP services, using industry standard SANS data on threats to various components of a firm’s information system over the period January 2003 to February 2011. Our results reveal strong evidence of contagion between such attacks, with attacks on ssh and Secure Web Server indicating increased attack activity on other ports. Security managers who ignore such contagious inter-relationships may underestimate the underlying risk to their systems’ defence of security attributes, such as sensitivity and criticality, and thus delay appropriate information security investments.

Keywords: Information Security, Security Attacks, Contagion, Hawkes Process, Security

Management

1. Introduction

Information systems and their hosted applications are typically subject to information security flaws, which, if exploited, may lead to substantial losses to the organization. When such threats

*Corresponding Author

appear, security managers will deploy mitigating responses to secure their systems.

A convenient classification for information security is determined by the concepts of *confidentiality*, *integrity*, and *availability*, or ‘CIA’. Informally, confidentiality is the property that just the right agents have access to specified information or systems, integrity is the property that specified information or systems are as they should be, and availability is the property that specified information or systems can be accessed or used when required. Convenient alternatives to confidentiality, integrity, and availability are *sensitivity* and *criticality*, in which sensitivity amounts to confidentiality together with some aspects of integrity and criticality amounts to availability together with some aspects of integrity. Here we characterize the security status of systems in terms of measures of their levels of sensitivity and criticality. Of course, it is possible to extend the vector of classifications to a larger number of attributes, at a cost of complexity, but the CIA-based view is generally understood to be comprehensive. Indeed, the choices of sensitivity and criticality fall out of the impact calculators provided by information sharing agencies such as the US National Institute for Standards in Technology (NIST) National Vulnerability Database (NVD).¹ The Common Vulnerability Scoring System (CVSS) within the NVD provides a calculator that firms can use to assess the impact of information security threats on their corporate networks. The calculation of a firm-specific impact is a function of the published vulnerability metrics and firm-specific characteristics. For example, the number of machines running a particular piece of vulnerable software, or the types of communications ‘ports’ that a networked computer relies upon, and how often these are scanned and attacked. Most Chief Information Security Officers (CISOs) will operate on some type of decision-rule that takes a large number numerical measurements on the characteristics of the threat and reduces them down to one or two measures of interest. Subsequently, using well-established definitions (see, for example, University of Georgia, Office of Information Security (2012) for a simple and elegant discussion), by *sensitivity*, we mean the level of security required for protecting data from access by unauthorized agents; by *criticality*, we mean the importance of the availability of accurate information for continuing system operations. Given that the time-variation in these two metrics is a function of both the threat and the firm specific characteristics, the realization of sensitivity and criticality will be idiosyncratic for each firm. Our results provide substantial empirical support for the theoretical notion that two combined metrics provide a comprehensive summary to the CISO in deciding on the reaction of their security posture to new events.

¹See <http://nvd.nist.gov>

We assume that threats to the security of systems represent potential losses to the operations of an organization. Security managers, in assessing the potential risks to an organization, are faced with a great deal of data concerning the distribution of attacks against the services provided by their systems. Intuitively, managers understand that there are significant relationships between attacks against different services: for example, compromising one service may enable an attack against another. The nature of the inter-relationship between the threats, if revealed, provides additional information to assist managers in making their choices of mitigating responses. For example, if the inter-relationship between threats is constant, independently of the frequency and intensity of threats, security managers can adopt smooth mitigation profiles to meet the threats. In the absence of such stable relationships, the managers' responses must be adjusted dynamically: for given temporal relationships between the number of attacks, their change (or 'jump') in frequency, and their change in size (extent of impact).

Our contribution to research on security effectiveness is to reveal an otherwise hidden aspect of the relationship between attacks. Using a model based on the mutually self-exciting Hawkes process, Hawkes (1970, 1971b,a); Aït-Sahalia et al. (2010), we demonstrate the presence of contagion between attacks against the different critical services — such as email, databases, name and directory servers, website operations, and shared storage — provided by the organization's systems. Our model parametrizes contagion, so allowing us to test the hypothesis that attacks are transmitted from one port to another. (Ports are the service- or process-specific software constructed to serve as communications endpoints in a computer's operating system.) We aim to assess the existence of contagious behaviour between these threats using threat data obtained by DShield and published by SANS (<http://feeds.dshield.org>). To our knowledge, our approach is new to the literature on the statistical structure of attacks.

Our technical contribution is threefold. First, we substantially extend the theoretical insights we developed in Ioannidis et al. (2012a) to include point processes with non-deterministic jump intensities and with an $n > 2$ -variate aggregation of realizations of attacks on a secure system.

The extension is essentially to re-cast the entire static decision process in Ioannidis et al. (2012a) into a framework where the vector of state variables, representing the risks to the network, evolves via a random process characterised by jumps with a time varying arrival rate. We then use this data-driven process to simulate a real firm's (albeit anonymized) reaction to 'port-scanning' threats.² Second, we

²Port scanning is a technique whereby an attacker probes ports, access points, on a network. Early port scanning looked for open ports to access a part of the network; however, modern techniques involve actively probing for out-of-date port protection to exploit vulnerabilities in closed or encrypted ports.

use the new derivation of the equilibrium decision function to illustrate the different time profiles of investment, with a variety of stochastic processes driving attacks. We then fit to a sample of attack data a general vector model that permits both continuous diffusions and contagion via two types of jumps. Third, we identify the critical components of the system for these attacks using an eigenvalue analysis of the contagion matrix.

Theoretical aspects of contagion in information security have been addressed using game theory in Parachuri et al. (2007); Lelarge and Bolot (2008); Lelarge (2009); Grossklags et al. (2008); Bachrach et al. (2011). These studies refer to the optimality of actions of both attackers and defenders and diverse system architectures. To our knowledge, there is no published empirical evidence regarding the statistical behaviour and inter-relationships between reported attacks on specific ports. Other, indirectly related work includes, for example, Böhme and Kataria (2006b,a); Böhme and Schwartz (2010), where citations to other background work can also be found. Why might attacking intensity be clustered or ‘lumpy’ in its intensity? The answer lies in the fixed costs attackers incur when developing tools, such as port scanners: as new techniques for exploiting vulnerabilities develop, there is aggregation of techniques to exploit these vulnerabilities over time and this aggregation results, as new tools with collections of exploits are brought online, in clustering of attacks.

The remainder of this paper is organized as follows: in Section 2, we outline a basic causal model that relates the existence of threats to the sensitivity and criticality security attributes of a system. The model focusses on the inter-relationship between the security attributes and the threats. In Section 3, we provide a characterization of the statistical methodology that is based upon the Hawkes process, Hawkes (1970, 1971b,a); Aït-Sahalia et al. (2010). The Hawkes process is a model of contagion, Aït-Sahalia et al. (2010). Section 4 discusses the data. The results and their implications for information security are presented in Section 5 and Section 6 presents a simulated example of a security manager’s choices, when investing to protect criticality and sensitivity. Section 7 presents our conclusions concerning the nature of the relationships between the threats and the differential informational significance of some threats.

2. The Basic Model

We consider a security manager who must trade off criticality (C), sensitivity (S), and investment (K). Deviations of criticality C_t and sensitivity S_t (as functions of time, t) from their long-run targets \bar{C} and \bar{S} , respectively, are linear functions of attacks on the various technological components of the

system represented by the random m -vector X_t . Where,

$$\{C_t - \bar{C}, S_t - \bar{S}\} = \{w'_C X_t, w'_S X_t\} \quad (1)$$

and w_C and w_S are m vectors of weights representing the vulnerability of the system to attacks (and $(\cdot)'$ denotes transpose). Under, the circumstances described above, the security manager has just one vector stochastic integral to evaluate,

$$X(t, T) = \int_t^T a(X_\omega | \theta) d\omega \quad (2)$$

where $a(\cdot)$ characterizes a multivariate càdlàg (continue à droite, limite à gauche) process see for instance Billingsley (1995) or Protter (2004) for the arrivals of attacks on the system, with vector of parameters θ .

The relative values of these weights (w_C and w_S) are determined, for a given instance of the model, by the relative intensities of attacks against the ports that are significant for these attributes. For the policy planner, the weights are assumed to be constant over a planning horizon $t < T$. The security manager, in effect, trades expected loss of criticality and sensitivity against deterministic investment expenditure.

For the purposes of risk management, the security manager trades off loss from criticality and sensitivity attacks against a costly additional investment over the planning horizon t, T . This is defined by the current level of investment K_t exceeding its long run target \bar{K} . Integrating over the time horizon, the appropriate loss functions from attacks and additional investment are given by, \bar{K} , with the appropriate loss functions given by

$$L_{CS}(t, T) = \int_t^T e^{-rt} f_C(C_t - \bar{C}) d\omega + \int_t^T e^{-rt} f_S(S_t - \bar{S}) d\omega \quad (3)$$

$$+ \int_t^T e^{-rt} f_{CS}(C_t - \bar{C}, S_t - \bar{S}) d\omega$$

$$L_K(t, T) = \int_t^T e^{-rt} f_K(K_t - \bar{K}) d\omega \quad (4)$$

where r is the discount rate and for simulation purposes is set to zero, ω is the sample space of

outcomes, and f_C , f_S and f_{CS} are affine functions that scale the measured deviations from target, $C_t - \bar{C}$ and $S_t - \bar{S}$, to a loss deriving from the deviation, $K_t - \bar{K}$, from the target profile of investment. The critical tipping point for additional investment occurs when the loss from additional investment ($L_K(t, T)$) equals the loss of criticality and sensitivity due to attacks ($L_{CS}(t, T)$). Similar models and conditions have been given in Ioannidis et al. (2009, 2012b,a). At this tipping point, the cost of inaction is equal to the cost of action (i.e. additional investment in information security) and this constitutes an optimality condition for the information security management.

We assume that \bar{K} , w_C , and w_S are determined by the policy and technology mix of the firm. For example, \bar{K} is allocated to the manager by the board and w_C and w_S are conditional on the systems architecture. This configuration is a choice made outside of the operational security planning phase. For example, a firm has a variety of operational choices for networks and client software. These will depend on the area of business of the firm, its size, and its exposure to attacks on its confidential information. Such configuration will determine \bar{K} , w_C , and w_S for each firm.

Following Ioannidis et al. (2012a), building on Ioannidis et al. (2012b, 2009), Equation 3 can be characterized uniquely by the stochastic process driving threats:

$$L_{CS}(t, T) = f_C(w'_C X(t, T)) + f_S(w'_S X(t, T)) + f_{CS}(w'_C X(t, T) X'(t, T) w_S) \quad (5)$$

That is, a second-order Taylor expansion of the loss functions. In Section 3 we introduce a process to describe the evolution of the vector X_t . This process is assumed to have two components, first a standard two dimensional Brownian motion and second a jump process, with normally distributed jumps. For the time horizon t, T the realisation for the statistical model used by the manager will therefore be a mixture of normals. If the unconditional variance, over (t, T) , of the jump process exceeds the variance of the continuous diffusion the resulting distribution of realisations (often referred to as the transition density function over t, T) will be symmetric and Leptokurtic.

Under the above distributional assumptions for X_t , we can filter for the higher moments by adjusting the unconditional second moments in the managers decision function without direct reference to the higher order moments in the actual data. This approach is taken in Section 3. Terms beyond the fourth moment are too small to consider for most regular distributions.

Several important aspects emerge from this decomposition. If the threats are independent diffusions with homogenous moments and co-moments, then security investment may be approximated by a ‘smooth investment profile’. By a smooth investment profile is understood that the security manager

knows the moments of these diffusions for some finite time horizon and plans their investment profile relative to their level of risk aversion as he knows the probability of exceedance accurately, therefore no discrete jumps are anticipated.

However, two further profiles may also exist: first, a set of independent, but self-exciting, point processes will characterize a set of time homogenous discontinuities in this investment horizon; second, mutually self-exciting jumps in the attack process will characterize highly localized discontinuities in the investment profile.

By ‘self-excitement’ we mean that the probability of an additional attack is conditional on the existence of a previous jump in attacks. By ‘mutual self-excitement’ the concept of self-excitement is extended by having jumps across other ports influencing the contemporaneous and future probability of jumps in a specific port. The magnitude of the relationship between ports is described by a time varying matrix containing all of the pairwise interactions of the individual elements of the vector X_t .

The investment choices made by security managers will be determined by the form of the loss function (Equation 5). The inter-temporal behaviour of the loss function depends on the behaviour of the attack process (Equation 2). We will model the attack process using a geometric Brownian motion (\mathcal{GBM}) with finite activity Poisson jumps and stochastic volatility. In general, this type of approach can capture many aggregate effects in the data generating process, for instance extreme events, high levels of auto-correlation and hence persistence in both mean and variance of attacking activity and most importantly temporary, but persistently high levels of covariation.

We can expand Equation 2 to display its contributing stochastic processes over the sample space s :

$$X(t, T) = X_0 + \underbrace{\int_t^T \mu_s ds}_{\text{drift}} + \underbrace{\int_t^T \sigma_s dW_s}_{\text{continuous part}} + \text{JUMPS} \quad (6)$$

$$\text{JUMPS} = \underbrace{\int_t^T \int_{\{|x| \leq 1\}} x(\mu - v)(ds, dx)}_{\text{finite activity jumps}} + \underbrace{\int_t^T \int_{\{|x| > 1\}} x\mu(ds, dx)}_{\text{infinite activity}} \quad (7)$$

where μ_s is the deterministic drift, σ_s is a (possibly stochastic) matrix volatility function, and JUMPS is the discontinuous point process, which is decomposed into finite and infinite activity terms. Here

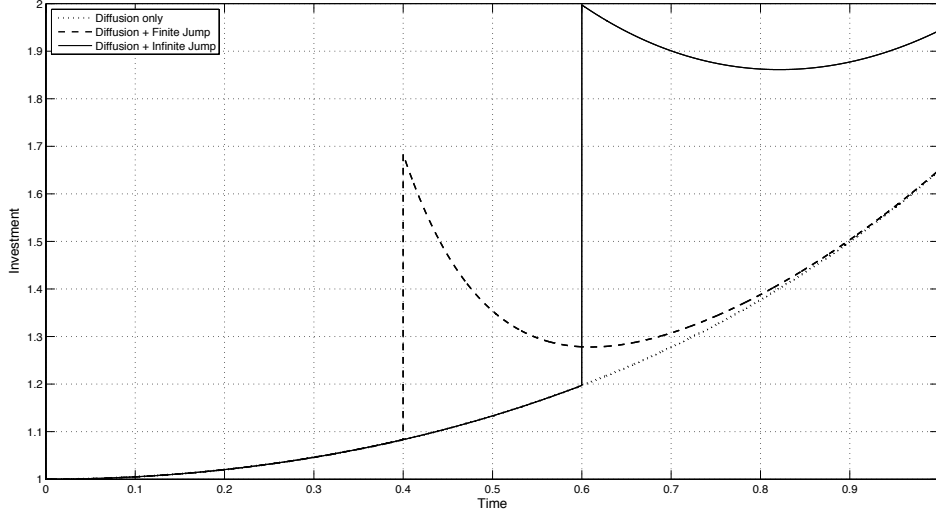


Figure 1: Illustration of various investment profiles, with and without jumps. The long run profile is based around attacks that are in the form of a continuous diffusion. Jumps are in two categories: jumps with finite activity, whereby the effect of the jump on the investment profile returns to the long run profile. Jumps with infinite activity never return to the long run profile, but form a new profile parallel to the long run.

X_0 is the initial endowment of X_t and $X(t, T)$ is the transition from t to $T > t$, usually assumed to be some regular planning horizon, such as a month or a quarter.

Finite activity jumps exhibit mean reversion in their intensities, where μ is a vector of average jump activity and ν is the rate of reversion to the mean level of jump activity. In contrast, infinite activity jumps are characterized by a vector of starting rates μ , but don't mean revert.

When $L_K(t, T) = L_{CS}(t, T)$, the dynamics of the attack effort determine the temporal profile of investment. The three contributing stochastic processes influence the investment profile as indicated in Figure 1. We revisit the issue of investment profiles, based on $L_K(t, T) = L_{CS}(t, T)$, in Section 6.

There is a variety of choices for the attack process X_t . These choices include doubly stochastic poisson point processes, with mean and non-mean reverting jump intensities, and systems of univariate self-exciting point processes, among others. Our choice is a multivariate form of a self-exciting point process, commonly referred to as a Hawkes process, that exhibits stochastic volatility and mutual self-excitation across the elements of the attack vector. This process was introduced in Aït-Sahalia et al. (2010).

3. The Statistical Methodology Based on the Hawkes Process

Vector processes that characterize attacks should allow for the following properties:

- Be represented as a vector stochastic differential equation (or, equivalently, stochastic integral equation) in order to support dynamic programming for forward-looking simulations and hence policy planning;
- Exhibit stochastic volatility in the diffusion process in order to capture potential changes in the variance of attacker behaviour and technological mechanisms of attacks;
- Exhibit discontinuous jumps that cluster in order to capture sudden changes in the vulnerability profile of systems/ports;
- Admit jumps that cluster across systems/ports, as attacks will most likely be multi-faceted (contagious).

The mathematical development in the remainder of this section establishes the necessary statistical methodology to support these features of our approach. The reader not wishing to follow the detailed mathematical development might proceed directly to Section 4.

We propose a statistical model describing the evolution of attacks to different ports over time based on the geometric Brownian motion with mean reverting stochastic volatility. This process is supplemented by a Poisson jump process to capture the possible discrete movements characterizing attacks to ports over some epochs.

We therefore decompose the integral from Equation 2 into mutually self-exciting jump diffusion processes, to capture potential contagion effects between attacks on different ports. Specifically, we adopt the Hawkes process Hawkes (1970, 1971b,a); Aït-Sahalia et al. (2010), a very general specification that captures the probability of jumps and allows for the parametric estimation of mutual self-exciting processes.

Equation 8 denotes the time evolution attacks (generalized in the integral Equation 6) to the i^{th} port as a stochastic differential equation. We assume that this process has three major components a deterministic drift term ($u_i dt$), a continuous variance term ($V_{i,t}$), and a jump term, dN of size Z .

$$dX_{i,t}/X_{i,t} = u_i dt + \sqrt{V_{i,t}} dW_{i,t}^X + Z_{i,t} dN_{i,t} \quad (8)$$

where $dW_{i,t}^X$ is a Brownian motion. The variance rate equation (9) is given a stationary stochastic process:

$$dV_{i,t} = k_i (\theta_i - V_{i,t}) dt + \eta_i \sqrt{V_{i,t}} dW_{i,t}^V \quad (9)$$

where $dW_{i,t}^V$ is a Wiener process, θ_i denotes the long-term variance, k_i the speed of adjustment, and η_i denotes the kurtosis. For the vector of Wiener processes the cross variation is denoted $R(t, T) =$

$\langle W_i, W_j \rangle_t^T$, with infrequent jumps, this will tend to the, scaled, long run unconditional correlation of the various processes. Alternatively, one can use a jump robust measure of cross variation (for instance the flat-top Kernel measurement cross variation of Barndorff-Nielsen et al. (2011)). Each port process is a jump augmented geometric brownian with square root variance process similar to those found in option pricing models, see Hull (2006) for alternative examples.

The jump process dN is assumed to be a Hawkes process, whose evolution can be expressed in terms of its intensity process $\lambda_{i,t}$,

$$\begin{cases} \mathbb{P}[N_{i,t+\Delta} - N_{i,t} = 0 | F_t] = 1 - \lambda_{i,t}\Delta + o(\Delta) \\ \mathbb{P}[N_{i,t+\Delta} - N_{i,t} = 1 | F_t] = \lambda_{i,t}\Delta + o(\Delta) \\ \mathbb{P}[N_{i,t+\Delta} - N_{i,t} > 1 | F_t] = o(\Delta) \end{cases} \quad (10)$$

where $N_{i,t+\Delta}$ is an m point process counting the number of jumps in $(0, t + \Delta)$ for the $i = 1, \dots, m$ processes in the system and $F_{i,t}$ is the conditional mean jump rate per unit of time. The jump intensities exhibit clustering according to the following dynamics:

$$\lambda_{i,t} = \lambda_{i,\infty} + \sum_{j=1}^m \int_{-\infty}^t g_{i,j}(t-s) dN_{j,s} \quad (11)$$

where $i = 1, \dots, m$ and $s \leq t$, and $j = 1, \dots, m$; the distribution of jumps $N_{j,s}$ is determined by that of the intensities $\lambda_{i,t}$, where $\lambda_{i,\infty}$ is the long-term intensity. Alternatively, Equation 11 can be expressed in terms of the integral over the sample space of outcomes ω . Therefore, we have an equivalency between the deterministic intensity adjustment for a known history of jumps over $t - s$ and the sample space ω , this is expressed as follows

$$\lambda_i = \lambda_{i,\infty} + \sum_{j=1}^m \lambda_j \int_{-\infty}^t g_{i,j}(t-s) ds = \lambda_{i,\infty} + \sum_{j=1}^m \left(\int_0^\infty g_{i,j}(\omega) d\omega \right) \lambda_j; \quad (12)$$

the vector function g is assumed to follow an exponential decay of the form

$$g_{i,j}(t-s) = \beta_{i,j} e^{-\alpha_i(t-s)} \quad (13)$$

for coefficients $\beta_{i,j}$ giving expected instantaneous jump values and decay rates α_i . In matrix form,

$$\Lambda_t = \Lambda_\infty + \Gamma_t \quad (14)$$

where Λ_∞ is an $m \times m$ diagonal matrix, whose diagonal elements are given by $\lambda_{i,\infty}$ and Γ_t is an $m \times m$ matrix, whose elements are

$$\gamma_{i,t} = \sum_{j=1}^m \lambda_j \int_{-\infty}^t g_{i,j}(t-s) ds = \sum_{j=1}^m \left(\int_0^\infty g_{i,j}(\omega) d\omega \right) \lambda_j \quad (15)$$

The overall association between the jumps of the different attacks is then captured by the matrix $G(\tau)$, where $\tau = t - s$ (following Ait-Sahalia et al. (2010), we take to be τ one day):

$$G(\tau) = \begin{pmatrix} \beta_{11}e^{-\alpha_1\tau} & \dots & \beta_{1m}e^{-\alpha_1\tau} \\ \vdots & \ddots & \vdots \\ \beta_{m1}e^{-\alpha_m\tau} & \dots & \beta_{mm}e^{-\alpha_m\tau} \end{pmatrix} \quad (16)$$

The diagonal elements indicate the self-excitation of the process — that is, when a jump occurs, the likelihood of another jump increases — whereas the off-diagonal elements indicate the influence of jumps in other attacks on the evolution of its own jumps. The existence of non-zero off-diagonal elements is indicative of the need, as discussed in the introduction, for managers to adjust their security responses dynamically according to their observations of the total threat environment.

If $G(\tau)$ is of full rank, then all of the elements of the attack process X mutually excite in the components of JUMP. In this case, all of the eigenvalues of $G(\tau)$ will be non-zero, so that every component of the vector λ_t contributes mutual self-excitation to every other component, with relative weights characterized by the eigenvectors.³

In our model, we are concerned with attack vectors against a range of ports. We would, therefore, expect that not every component will contribute significantly to the mutual self-excitation of the whole system: some will simply be of little systemic relevance. Consequently, with some eigenvalues approaching zero, $G(\tau)$ can, in practice, be treated as being of reduced rank.

Dividing each eigenvalue by the sum of the eigenvalues (the trace of the diagonal matrix) gives

³In general, $G(\tau)$ need not be positive semi-definitive, because of the existence of possible asymmetric responses in the intensity process. In the case that $G(\tau)$ is positive semi-definite then the resulting eigenvectors are the weights of a set of a set of orthogonal processes that are interpretable as principal components.

the proportionate contribution of each component to the total mutual self-excitation of λ_t . Now the eigenvectors of the largest eigenvalues will determine the key features of the model.

For the system of vector equations,

$$\begin{aligned} dX_t &= udt + \sqrt{V_t}dW_t^X + Z_t dN_t \\ dV_t &= \kappa(\theta - V_t) dt + \eta\sqrt{V_t}dW_t^V \\ d\lambda_t &= \alpha(\lambda_\infty - \lambda_t) dt + \beta dN_t \end{aligned}$$

Aït-Sahalia et al. (2010) identify the first two moment conditions as the expectations

$$\begin{aligned} \mathbb{E}[\Delta X_t] &= (\mu + \lambda M[1]) \Delta + o(\Delta^2) \\ \mathbb{E}[(\Delta X_t - \mathbb{E}[\Delta X_t])^2] &= (\theta + \lambda M[2]) \Delta + \frac{\beta\lambda(2\alpha - \beta)}{2(\alpha - \beta)} M[1]^2 \Delta^2 + o(\Delta^2) \end{aligned} \quad (17)$$

For identification of all the parameters, two additional moment conditions (skewness and kurtosis) are required. These are given in Equations 18 and 19.

$$\begin{aligned} \mathbb{E}[(\Delta X_t - \mathbb{E}[\Delta X_t])^3] &= \lambda M[3] \Delta \\ &+ \frac{3}{2} \left(\eta\theta\rho^V + \frac{(2\alpha - \beta)\beta\lambda M[1] M[2]}{(\alpha - \beta)} \right) \Delta^2 + o(\Delta^2) \end{aligned} \quad (18)$$

where ρ^V is the first-order autocorrelation coefficient of the intensities $\lambda_{i,t}$ and the $M[i]$ indicate the centred moments matrices, and the fourth moment (kurtosis) as

$$\mathbb{E}[(\Delta X_t - \mathbb{E}[\Delta X_t])^4] = \lambda M[4] \Delta \left(\begin{array}{l} \frac{3\theta\eta^2}{2\kappa} + 3\theta^2 + 6\theta\lambda M[2] + \\ 3\lambda \left(\lambda + \frac{(2\alpha - \beta)\beta}{2(\alpha - \beta)} \right) M[2]^2 + \\ \frac{2(2\alpha - \beta)\beta\lambda M[1] M[3]}{(\alpha - \beta)} \end{array} \right) \Delta^2 + o(\Delta^2) \quad (19)$$

The parameters of the model in Equations 17 and 19 — $\mu, \theta, \beta, \lambda, \alpha, \eta, \rho^V$, capturing the possible correlations between the Brownian motions and κ — are estimated by finding the roots of the system given by Equations 17 and 19. For more details of the estimation procedure via Generalized Method of Moments (GMM), see Aït-Sahalia et al. (2010).

The Matlab code provided by Aït-Sahalia et al. (2010) utilizes the Mathworks optimization toolbox to minimize the theoretical and empirical moments. The Hessian function at the minima provides

the standard errors for the parameters for use in statistical inference.

The predicted values from the first two moment conditions in Equation 17 can be substituted into Equation 5 enabling us to calculate the expected losses due to the flow of attacks. On the basis of this methodology we compute the moments of the number of attacks on different ports as a proxy for the possible losses in criticality and sensitivity $L_{CS}(t, T)$ over a planning horizon t, T . Subsequently to the statistical analysis we proceed in simulating potential investment paths based on the statistical predictions obtained above.

4. Data

Getting a reliable picture of the attack environment is very difficult. Most organizations are very sensitive about the details of a attacks and how they happen. Such information is rarely shared. This makes it very hard to understand the attack environment and to estimate how it will evolve as, for example, the use of cloud increases.

Organizations have many sources of information about attacks that may be incident upon their networks. One source of particular interest is firewall logs. Most, if not all, corporate networks will run a firewall that limits the traffic in and out of the corporate intranet according to some set of rules. Firewalls also log the network activity that they see, particularly the network traffic that is being dropped. Security teams examine firewall logs to get an indication of what attacks are occurring. The log files may show particular IP addresses that are running scans or particular network ports that are being attacked.

The dataset upon which this paper draws is taken from the output of the DShield community project (<http://feeds.dshield.org>).

- DShield is a community project — sponsored by SANS (<http://www.sans.org>) — that correlates firewall log files from many volunteer companies in order to paint a picture of the current threat environment.
- DShield consists in a client system that converts firewall log files (from many different vendors) into a standard format. These are then sent to a data collection engine where the data is aggregated and used to analyze attack trends.
- DShield has very wide global coverage and has become the dominant attack correlation engine. It has been used in the early detection of worms and is used to analyse attack patterns. Much of

the work using DShield data has been done to analyze particular events rather than to understand the overall attack space.

- DShield has been collecting data for close to a decade so in this paper we take a historical view of the data looking at its statistical properties rather than individual events.

For our statistical analysis, we picked ten particular services, sampled daily for the period 1 January 2003 to 28 February 2011. The data was extracted from the SANS DScale database on 1 March, 2011. Data for each of the ports of interest was collected. For example, for port 53, <https://isc.sans.edu/portascii.html?port=53&start=2003-01-01&end=2011-02-28>. The data was processed to fill in missing dates, with missing values filled using piecewise cubic spline interpolation. The number of missing points in our sample lies between 1/2 and 1% of the total observations for each series.

The services considered are given in Table 1 and the descriptive statistics are given in Table 2. Here we looked for ports that would typically be used to run services that we would expect to see offered as part of a cloud service (either to manage the service or to help build other services). We also included DNS as its correct functioning is fundamental to the internet. The ports considered are significant for different security attributes, such as sensitivity and criticality, to varying relative extents. For example, the Secure Web Server is highly significant for sensitivity and DNS is highly significant for criticality.

Services such as ssh and DNS are enablers for most of the operations of the internet and as such are highly attractive targets for attackers wishing to interrupt other services. As such they will be subject to relatively constant high levels of attacks (indicated by high mean, low variance and low kurtosis).

However, applications such as oracle do not enable other services to the same extent. Oracle, Secure Web Server and IMAP would generally be expected to have high levels of variation in implementation. For this type of measurement a greater diversity of implementation would result in higher variation in security vulnerabilities (indicated by a relatively low mean, high variance and very high Kurtosis).

Since the primary purpose of this paper is to illustrate a model, we have not attempted to filter the DShield data for false positives.

Table 1: Services considered in extracts of DShield attack data (<http://feeds.dshield.org>)

Service	Port Number	Description
DNS	53	A service used to find the IP address of a particular service given its name
ssh	22	Secure shell. A program used to connect to computers remotely
Oracle	80, 443	A popular enterprise database used at the core of many business applications
SQL	118	Microsoft's database which is again used at the heart of many business applications
LDAP	389	A directory service that often contains the name and details of employees within a company and which is used to determine employees' rights to access business applications
Web Server	80	Used to run websites. There are many different applications that could be used here but popular ones are IIS and Apache
Secure Web Server	443	The secure part of a web server where traffic is encrypted using SSL. Usually used for highly sensitive transactions
Samba	139, 455	A shared drive used to store and share information within many organizations
Email (IMAP)	143, 993	The protocol used by many email clients to access an email server. Many web based email services also support this protocol
Email (SMTP)	25, 465	SMTP is used by some email clients to send an email to an email server, but it is also used to forward emails between different email servers as email is sent from the sender's email server to the recipient's

Table 2: Daily Time Scale Descriptive Statistics of the Attack Series.

	DNS	ssh	Oracle	SQL	LDAP
Mean	331529	2782020	8334	929074	43042
Standard Deviation	307232	2964248	46118	867475	47517
Skewness	4	2	17	6	2
Kurtosis	33	11	345	67	7

	Web Server	Secure Web Server	Samba	IMAP	SMTP
Mean	787514	69568	200394	1426	20235
Standard Deviation	998963	83628	207624	4468	45661
Skewness	4	11	6	15	5
Kurtosis	22	250	72	329	36

5. Results and Analysis

Following the statistical methodology outlined in Section 3, we have estimated the vector equation system given in Equations 17, 18, and 19 by GMM. Inference was undertaken using the estimated information matrix of the system. Given the challenging dimensionality of the system, we are reporting a selection of our statistical results, which shed light on our choice of statistical model. Our model incorporates stochastic jumps to the geometric Brownian motion describing the evolution of attacks and also allows for the existence of mutually self-exciting processes. The parameters of interest are, therefore, represented by the estimates of λ_∞ that are indicative of the existence of jumps. Contagion is captured by the elements of the contagion matrix (16).

All of the statistical results that are presented in the tables below are statistically significant at 5% and the original GMM information matrices are available on request. ⁴

Table 3 presents our estimates of λ_∞ and the diagonal elements of the contagion matrix (16) using expressions and , calculated from the estimated parameters derived from the optimization of (19). Our estimated long-run intensities, λ_∞ , are modest and have similar values across all time series of attacks. Their statistical significance indicates that our choice of the inclusion of jumps in the law of motion of attacks is justified and that jumps constitute a significant, albeit unobserved, component of attacks.

Figures 2 and 3 (below) depict, respectively, the number of attacks and the estimated jumps for the first five listed services. The graphs for the second five would be similar. Informally, the graphs suggest that the correlation between the jumps will be stronger than can be inferred from the raw data.

Our estimates of the parameters β and α are used as inputs to calculate the elements of the adjusted covariance matrix $G(\tau)$ and the second row of Table 3 reports our calculation of the diagonal elements of the contagion matrix based on these estimates which (16) provides strong evidence of the existence of self-exciting processes as both the diagonal and off-diagonal elements constituted in terms of the estimated parameters, are functions statistically significant parameter estimates. It is remarkable that Oracle and the Secure Web Server exhibit the highest degree of self-excitation followed by ssh and IMAP, this tallies with the technical properties of these two ports in providing active data and secure active web content. SSH is the secure shell that allows network access remotely and IMAP is an email hosting system. These four ports are the essential components of most business information systems, it appears that they are also the core features of any contagion in attacking behaviour. Compromises

⁴The data, routines and all the pivotal statistics for the parameter estimates are available from the authors webpage.

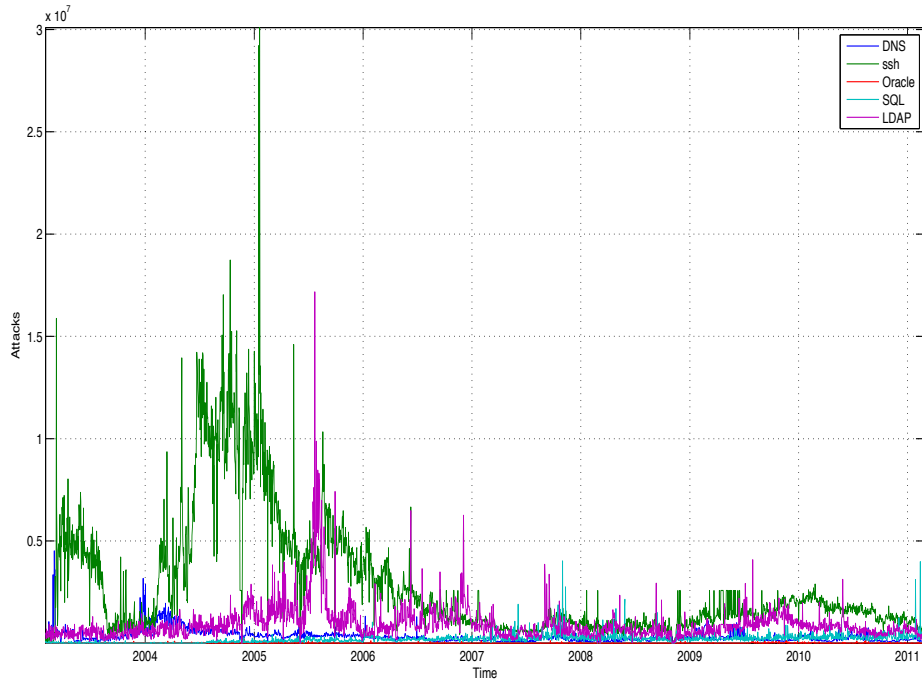


Figure 2: Daily attacks to five ports. For exposition purposes we plot five of the series in our data sample.

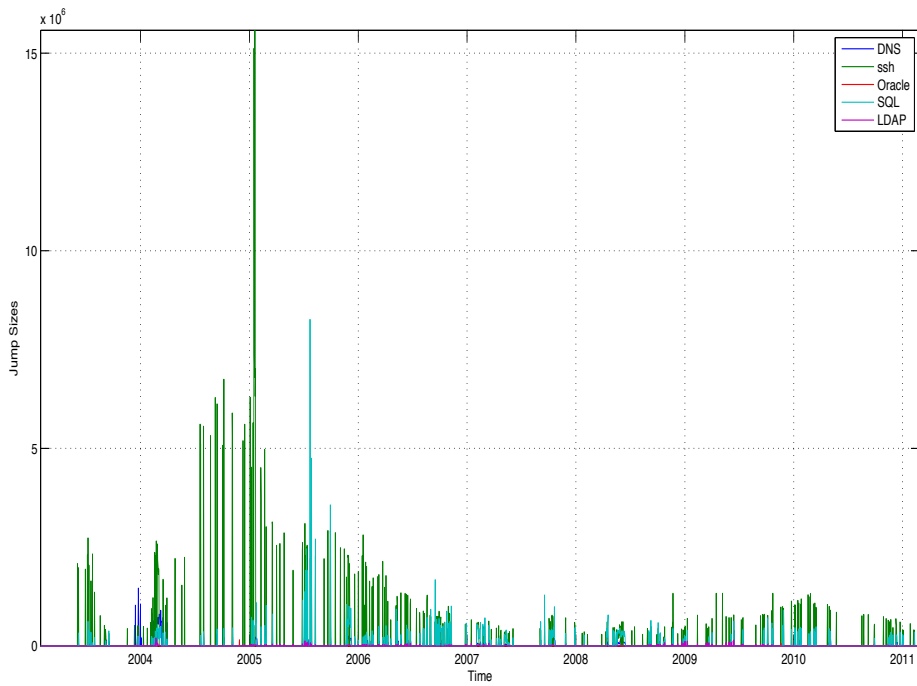


Figure 3: Extracted jumps for five ports. From the chosen ports in Figure 2 we report the extracted jumps from our decomposition.

(or the potential to compromise) in one or several of these systems appear to provide access (or the potential to access) that spills over to other systems very rapidly.

Table 4 presents our estimates of the normalized elements of the correlation (contagion) matrix. The unconditional correlation matrix contains the correlations of both the Weiner processes and the infrequent jumps, whilst the correlation (contagion) matrix only contains the conditional correlations between the excited jumps. (16). These two tables show large and very large difference between their elements. Whilst the unconditional correlations tend to be of modest magnitude, occasionally negative, whilst several are statistically insignificant, the elements occupying the same position in the normalised G-matrix are all large, positive and individually statistically significant at 1Based on the comparison of these estimates with the corresponding values in the unconditional correlation matrix given in Table 5, we provide strong evidence of the existence of uniformly high and positive correlations between attacks, so justifying our choice of the Hawkes process to capture evident of mutual self-excitement.

Once the jumps have been taken into account, the correlations between the intensity of attacks and their size are presented in Table 6. The structure of this matrix reveals that jumps cluster in both intensity and size, and that their association is almost uniform and very strong. This correlation structure indicates that time evolution of the set of attacks will exhibit periods of intense activity and large size of attacks, and other periods where such activity is very low.

Table 3: Long-run Intensities; Diagonal Elements of G

	DNS	ssh	Oracle	SQL	LDAP
λ_∞	0.1143	0.1158	0.1146	0.1114	0.1136
$\beta_{i,j}e^{-\alpha_i\tau}$	0.0714	0.0831	0.17	0.05	0.0632
	Web Server	Secure Web Server	Samba	IMAP	SMTP
λ_∞	0.1118	0.1125	0.1132	0.115	0.1125
$\beta_{i,j}e^{-\alpha_i\tau}$	0.0728	0.1463	0.0443	0.0928	0.0085

The eigenvalues of the estimated $G(\tau)$ matrix are 0.008, 0.0219, 0.0347, 0.0582, 0.084, 0.1328, 0.2664, 0.5403, 0.9894, and 7.8643. We have computed the corresponding eigenvectors of the estimated contagion matrix (Table 4). The two highest eigenvalues associated with the filtered (i.e., to include jumps and mutual self-excitation) time series are 0.9894 and 7.8643. These represent approximately 90% of the excitation. This establishes the differential information content of attacks on each

Table 4: Normalized G Matrix. Note WS refers to the 'Web Server' and SWS refers to the 'Secure Web Server'.

	DNS	ssh	Oracle	SQL	LDAP	WS	SWS	Samba	IMAP	SMTP
DNS	1	0.86	0.84	0.83	0.49	0.91	0.73	0.92	0.81	0.97
ssh	0.86	1	0.72	0.71	0.57	0.94	0.63	0.79	0.95	0.83
Oracle	0.84	0.72	1	0.99	0.41	0.76	0.88	0.91	0.68	0.86
SQL	0.83	0.71	0.99	1	0.41	0.75	0.89	0.89	0.67	0.85
LDAP	0.49	0.57	0.41	0.41	1	0.54	0.36	0.45	0.61	0.48
WS	0.91	0.94	0.76	0.75	0.54	1	0.67	0.84	0.89	0.88
SWS	0.73	0.63	0.88	0.89	0.36	0.67	1	0.79	0.59	0.75
Samba	0.92	0.79	0.91	0.89	0.45	0.84	0.79	1	0.75	0.95
IMAP	0.81	0.95	0.68	0.67	0.61	0.89	0.59	0.75	1	0.79
SMTP	0.97	0.83	0.86	0.85	0.48	0.88	0.75	0.95	0.79	1

Table 5: Unconditional Correlation Matrix.

	DNS	ssh	Oracle	SQL	LDAP	WS	SWS	Samba	IMAP	SMTP
DNS	1	0.46	0.02	-0.37	0.21	0.44	-0.01	0.09	-0.10	-0.26
ssh	0.46	1	0.18	-0.14	0.50	0.42	0.03	0.12	0.12	-0.26
Oracle	0.02	0.18	1	0.02	0.18	0.05	0.08	0.01	0.04	0.10
SQL	-0.37	-0.14	0.02	1	0.26	-0.37	0.17	0.38	0.51	0.47
LDAP	0.21	0.50	0.18	0.26	1	0.18	0.11	0.32	0.35	-0.06
WS	0.44	0.42	0.05	-0.37	0.18	1	0.15	-0.24	-0.23	-0.22
SWS	-0.01	0.03	0.08	0.17	0.11	0.15	1	-0.06	-0.08	0.26
Samba	0.09	0.12	0.01	0.38	0.32	-0.24	-0.06	1.00	0.45	0.01
IMAP	-0.10	0.12	0.04	0.51	0.35	-0.23	-0.08	0.45	1	0.04
SMTP	-0.26	-0.26	0.10	0.47	-0.06	-0.22	0.26	0.01	0.04	1

Table 6: Correlation Between Services.

	DNS	ssh	Oracle	SQL	LDAP	WS	SWS	Samba	IMAP	SMTP
DNS	1	0.86	0.86	0.89	0.88	0.81	0.84	0.87	0.89	0.87
ssh	0.86	1	0.86	0.86	0.83	0.83	0.84	0.84	0.87	0.87
Oracle	0.86	0.86	1	0.85	0.89	0.83	0.84	0.85	0.87	0.84
SQL	0.89	0.86	0.85	1	0.88	0.8	0.86	0.88	0.9	0.88
LDAP	0.88	0.83	0.89	0.88	1	0.86	0.83	0.87	0.9	0.83
WS	0.81	0.83	0.83	0.8	0.86	1	0.86	0.84	0.83	0.85
SWS	0.84	0.84	0.84	0.86	0.83	0.86	1	0.89	0.87	0.83
Samba	0.87	0.84	0.85	0.88	0.87	0.84	0.89	1	0.89	0.87
IMAP	0.89	0.87	0.87	0.9	0.9	0.83	0.87	0.89	1	0.86
SMTP	0.87	0.87	0.84	0.88	0.83	0.85	0.83	0.87	0.86	1

Table 7: Eigenvectors for the Two Highest Eigenvalues

DNS	ssh	Oracle	SQL	LDAP	Web Server	Secure Web Server	Samba	IMAP	SMTP
-0.0082	-0.0919	-0.4530	0.1689	-0.5110	-0.2563	0.3977	0.2114	-0.3618	0.3118
0.1218	0.7169	-0.0236	0.2244	-0.2364	0.1687	-0.4430	0.1447	0.0546	0.3403

port. From the corresponding eigenvectors, we infer that attacks on ssh and Secure Web Server are indicative of additional intensity and size of attacks on the remaining ports.

The totality of our statistical results supports our choice of filtering and reveals strong positive associations between attacks on the chosen ports, in contrast to the information suggested by the raw data that is indicative of weak associations, some of which have negative values. In the absence of our filtering, the use of the raw data might lead to erroneous responses by security managers in the face of attacks on specific ports. More specifically, the unconditional correlation between ssh attacks and IMAP is positive, but only by 12%, but for the filtered series the corresponding contagion matrix element $G(\tau)$ is positive and seven times larger at 0.87. In principle the ssh ports are important because of the associated access rights to the totality of the network. This is also important for the DNS and Secure Web Server ports, to a greater or lesser extent, depending on the specific set-up of the corporate network. The contagion effects implied within the decomposition appear to map to the conventional system architecture, we believe this to be the first result of its type.

6. A Simulated Investment Example

Consider a security manager with targets for deviations from criticality and sensitivity, as defined in (1). We now set up a security investment scenario in which losses to criticality and sensitivity are proportional to the attacks recorded in the SANS data set given in Section 4. Using the model given in Section 2, we recursively (one period forward at a time) estimate the investment manager's choices from 1 January, 2003 to 28 February, 2011. To do this we use an expanding window and assume single period myopia (therefore each individual iteration is a once for all independent decision).

For simplicity, we base our weights for the level of criticality w_C and sensitivity w_S on the first two eigenvectors of the long-run covariance matrix. Also for simplicity, we assume separable additivity; therefore, $f_{CS} = 0$.

We set f_C and f_S using the the approach in Ioannidis et al. (2012a). First, we assume that the underlying utility function is specified in terms of $u(-w_C X_t, -w_S X_t)$. We then follow Ioannidis et al. (2012a) by setting $u(-w_C X_t, -w_S X_t)$ to be from a family of hyperbolic utility functions with increasing absolute risk aversion (implying constant relative risk aversion). Assuming that γ_C and γ_S are relative risk-aversion coefficients, then the optimization required to attain a maximum utility will be the minimization of a loss function derived from the second order Taylor expansion of $u(-w_C X_t, -w_S X_t)$. Therefore,

$$f_C(w_C X_t) = E(w_C X_t) + \gamma_C E(w_C X_t)^2 \quad (20)$$

$$f_S(w_S X_t) = E(w_S X_t) + \gamma_S E(w_S X_t)^2 \quad (21)$$

For the purposes of this example, we set a discount rate of 3% per annum and assume that investment targets are set quarterly and last for one quarter. We assume that the policy manager has risk aversion in line with that of the military security planner discussed in Ioannidis et al. (2012b) and set $\gamma_C = 2$ and $\gamma_S = 4$.

Figure 4 plots the planned (continuous) and abnormal (dashed) investment for a firm facing attacks proportional to the intensity of port scans in the SANS data. The red investment line presents a planner setting forward investment from a brownian motion, whilst the black line presents a security planner using the jump diffusion model with mutually self-exciting jumps estimated previously.

Figure 4 presents a useful summary of the issues associated with not including an appropriate provision for jumps in the investment profile. The abnormal investment line (dashed) is the level of investment with perfect foresight. When the dashed line rises above the solid line, extra investment is required to mitigate a particular set of attacks.

The total loss for a security planner over the sample period is plotted in Figure 5. The log scale illustrates the substantial improvement in planning brought about by the inclusion of jumps. Mutual excitement allows the planner to allocate extra resources during periods of activity. In contrast, the rolling Brownian motion underestimates investment in periods of jump activity and overestimates investment in periods when jumps are absent.

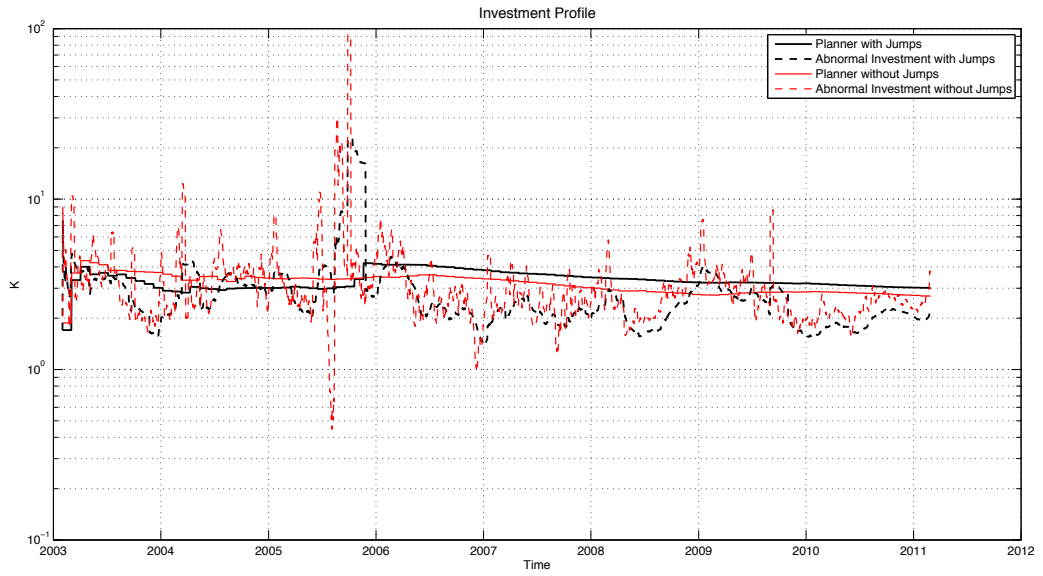


Figure 4: The investment profiles of a 'with jump' (black) versus 'without jumps' (red) security manager. Abnormal investment occurs when the dashed line is above the unbroken line.

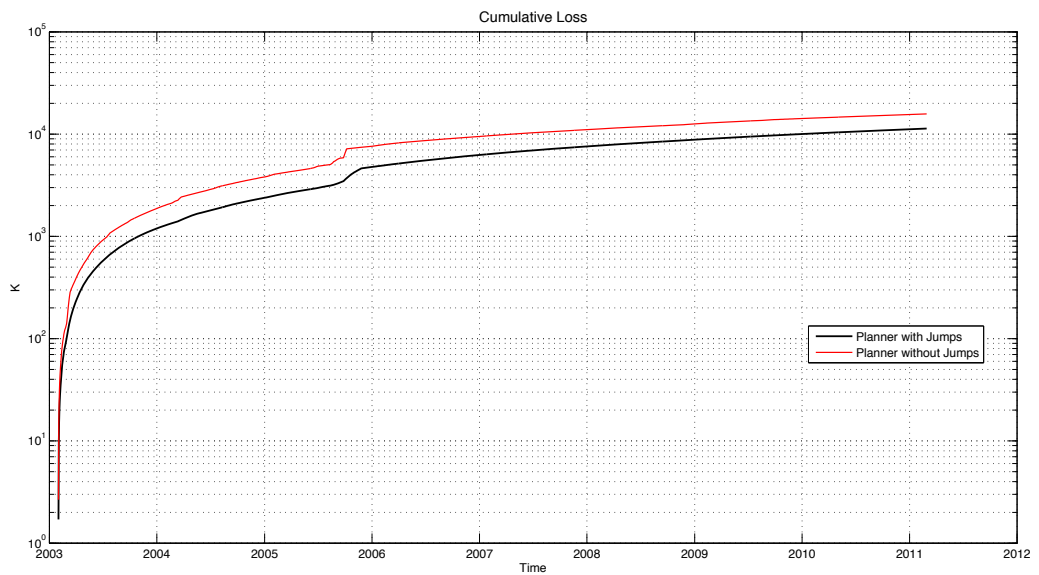


Figure 5: Cumulative total loss from 1 January, 2003 to 28 February, 2011 for a simulated security manager according Equation 3.

7. Conclusion

Our statistical analysis has revealed that attacks on individual ports are inter-related, with the relationships being exposed by the estimation of jumps and mutually self-exciting behaviour. From the ten chosen services, attacks on ssh and the Secure Web Server account for almost all of the self-exciting behaviour. Failure to reveal such relationships may lead to erroneous investment profiles. For example, the unconditional correlation between ssh (the compromise of which brings high levels of threat to both criticality and sensitivity) and the Secure Web Server (typically used to support transactions requiring a high level of sensitivity protection) is approximately zero; but, after filtering, it is 63% in the excited state. The responses of security managers should be based on this latter degree of correlation as it represents the impact of contagion on the level of risk that must be anticipated.

The efforts of security managers to protect attributes such as sensitivity and criticality at desired levels will necessitate additional costly investment at irregular intervals when faced with increased volume and diversity of attacks on, in particular, ssh and Secure Web Server. Such additional attacks will be associated with other additional attacks on all of the services considered.

We see from the data that the attack vector is characterized by mutually self-exciting finite and infinite jumps. Assuming that changes in attack intensity are related to changes in the rewards for attackers, then the loss function for the security manager will exhibit discontinuities, so that — if $L_K(t, T) = L_{CS}(t, T)$ — the corrective investment action will exhibit similar discontinuities. If the security manager knows this, then anticipatory investment can be made in advance of potential jumps by estimating times between jumps in the critical attack components.

This behaviour is in contrast to that which would be expected from a naïve security manager who assumes that the attack process is jump-free. Such a manager may, for example, set his investment profile to be an affine function of the drift, with an increment for risk-aversion proportional to σ_s , so that when jumps arrive he may suddenly be faced with unanticipated investment requirements.

In this case, the naïve manager plans from the unconditional moments of the attack process, thus adopting a smooth investment profile, based on these calculations. The failure to recognise the existence of jumps will result in over- investment in information security when no jumps are present and more importantly a substantial exposure to loss to criticality and sensitivity L_{CS} due to the inability to recognise the nature of the increase. This results in an underinvestment in security after a jump occurs, therefore total loss is always higher than the case when jumps are correctly included in the investment planning model.

An extension of the model would be to use the methodology in order to predict future levels of threat, so helping security managers to anticipate appropriate levels of investment. A refinement of the model would be to consider in more detail the mapping between the threats to services and the security attributes (such as sensitivity and criticality, or CIA) that may be compromised. Such an analysis might suggest how system architecture might be designed in order to limit the transmission of threats between services.

Other further work would be to consider how the behaviour of hackers correlates with our analysis of attacks. Such consideration is beyond the scope of this short paper.

Acknowledgements

We are grateful to Yacine Aït-Sahalia and Jean Jacod for the Matlab code used in the estimation procedure for the system described for Equations 19 to 17, the full derivation of the estimator is provided on pages 45 to 47 of Aït-Sahalia et al. (2010). The data and codes used in this paper is available from the authors' websites.

References

- Aït-Sahalia, Y., J. Cacho-Diaz, and R. J. Laeven (2010, March). Modeling financial contagion using mutually exciting jump processes. Working Paper 15850, National Bureau of Economic Research.
- Bachrach, Y., M. Draief, and S. Goyal (2011). Security games with contagion. Manuscript, 2011: <http://www.econ.cam.ac.uk/faculty/goyal/wp11/securitygames17.pdf>.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2011). Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics* 162(2), 149 – 169.
- Billingsley, P. (1995). *Probability and Measure*. John Wiley & Sons.
- Böhme, R. and G. Kataria (26–28 June, 2006b). Models and measures for correlation in cyber-insurance. In R. Anderson (Ed.), *Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS 2006)*, Robinson College, University of Cambridge, <http://weis2006.econinfosec.org>. <http://weis2006.econinfosec.org/docs/16.pdf>.
- Böhme, R. and G. Kataria (October 23–24, 2006a). A closer look at attack clustering. In S. Schecter (Ed.), *Proceedings of the I3P Workshop on the Economics of Securing the Information Infrastructure, Washington DC*, <http://wesii.econinfosec.org/workshop/>. http://wesii.econinfosec.org/draft.php?paper_id=35.
- Böhme, R. and G. Schwartz (June 7–8, 2010). Modeling cyber-insurance: Towards a unifying framework. In T. Moore (Ed.), *Proceedings of the Ninth Workshop on the Economics of Information Security (WEIS 2010)*, Harvard, <http://weis2010.econinfosec.org>. http://weis2010.econinfosec.org/papers/session5/weis2010_boehme.pdf.
- Grossklags, J., N. Christin, and J. Chuang (2008, June). Security investment (failures) in five economic environments: A comparison of homogeneous and heterogeneous user agents. In *Proceedings (online) of the Seventh Workshop on the Economics of Information Security (WEIS)*, Hanover, NH.
- Hawkes, A. (1970). Bunching in a semi-markov process. *J. Appl. Prob.* 7, 175–182.
- Hawkes, A. (1971a). Point spectra of some mutually exciting point processes. *J. Roy. Statist. Soc. B* 33, 438–443.
- Hawkes, A. (1971b). Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 83–90.
- Hull, J. C. (2006). *Options, Futures and Other Derivatives*. Prentice-Hall, London.
- Ioannidis, C., D. Pym, and J. Williams (2009). Investments and trade-offs in the economics of information security. In R. Dingledine and P. Golle (Eds.), *Proc. Financial Cryptography and Data Security '09*, Volume 5628

- of *LNCS*, pp. 148–166. Springer. Preprint available at <http://homepages.abdn.ac.uk/d.j.pym/pages/IoannidisPymWilliams-FC09.pdf>.
- Ioannidis, C., D. Pym, and J. Williams (2012a). Fixed costs, investment rigidities, and risk aversion in information security: A utility-theoretic approach. In B. Schneier (Ed.), *Economics of Security and Privacy III*, pp. 171–192. Springer.
- Ioannidis, C., D. Pym, and J. Williams (2012b). Information security trade-offs and optimal patching policies. *European Journal of Operational Research* 216(2), 434–444.
- Lelarge, M. (2009). Economics of malware: Epidemic risks model, network externalities and incentives. In *Communication, Control, and Computing*.
- Lelarge, M. and J. Bolot (2008). Network externalities and the the deployment of security features and protocols in the internet. In *SIGMETRICS*.
- Parachuri, P., J. Pearce, M. Tambe, F. Ordonez, and S. Kraus (2007). An efficient heuristic approach for security against multiple adversaries. In *AAMAS*.
- Protter, P. (2004). *Stochastic Integration and Differential Equations 2nd Ed.* Springer.
- University of Georgia, Office of Information Security (2012). Information Classification Standard. <http://infosec.uga.edu/policies/classification.php>.