Contagion in Cybersecurity Attacks

Berlin, June 2012

Adrian Baldwin, HP Labs, Bristol
Iffat Gheyas, University of Aberdeen
Christos Ioannidis, University of Bath
David Pym, University of Aberdeen
Julian Williams, University of Aberdeen

June 25, 2012

Talk

What we will cover:

- The idea behind the model and some prior studies in this area.
- How the model works and why it is a departure from prior models in this field.
- Our first set of results and the sample dataset of attack data.
- How to interpret them.
- Some conclusions and our future directions.

Introduction

Motivation

- This paper is part of an ongoing set of research projects in cyber and cloud security.
- Part of our work has been looking at the interaction between defensive expenditure and behavior versus attacker behavior and participation in threats.
- This paper is designed to look solely at the attack side and motivate some points regarding the clustering of cyber attacks.

Underlying Idea

- If attackers adjust their focus dynamically through time and across systems then we have prima facie evidence for the presence of attacker response functions.
- The key here is in the mutual and self excitation of vectors of attacks.

Related Literature

- 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, Draief, and Goyal (Bachrach et al.).
- These studies refer to the optimality of actions of both attackers and defenders and diverse system architectures.
- See for instance Böhme and Kataria (2006a,b); Böhme and Schwartz (2010), where other background work can also be found.
- Very recent work by the authors has looked at attack and defense problems when attackers choose to enter the market for attacks, based on expected reward versus expected costs.
- The dynamic equilibrium form of this model predicts attacks clustering, in time and across system attributes. This paper seeks to find evidence for this prediction.

Choice of Approach

- 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 m-vector X_t. Therefore

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

- where w_C and w_S are m vectors of weights representing the vulnerability of the system to attacks (and $(\cdot)'$ denotes transpose).
- For the policy planner, the weights are assumed to be constant over a planning horizon t, T.

The Attack Vector

- In previous papers we have looked at the dynamics of investment functions under a variety of threats.
- In this paper we shall look at the dynamics of the threats to systems and demonstrate the resultant shapes of investment functions, for this type of behaviour.
- These results are important, not only for our current research for industry policy makers, but for our forthcoming work on public policy.

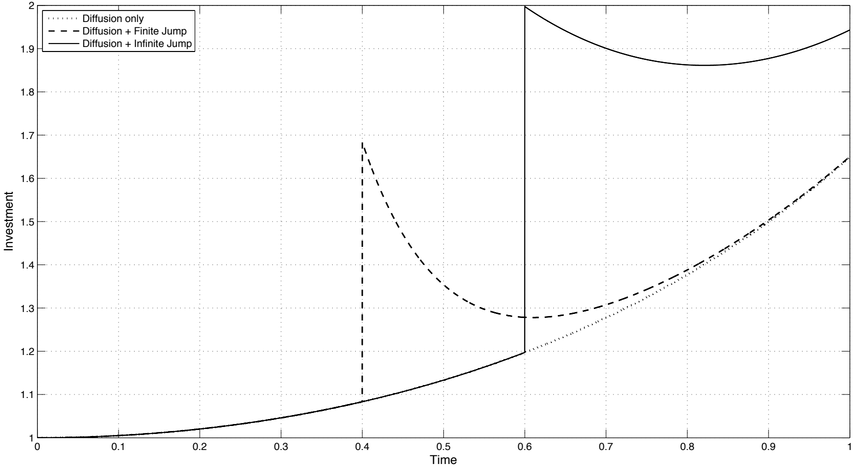
Diffusion only
Diffusion + Finite Jump
Diffusion + Infinite Jump
Investment
Time
2
1.9
1.8
1.7
1.6
1.5
1.4
1.3
1.2
1.1
1
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

The Model

Contagion Models

- Single equation models of self excitation date back to the 1970s, Hawkes (1970, 1971b,a); Aït-Sahalia et al. (2010).
- Multivariate models of mutual and self excitation are far more recent.
- Our model is based on the work by Aït-Sahalia et al. (2010) that generalizes the Hawkes process and identify the characteristic function and hence the GMM estimator for this very flexible process.
- This process admits the diffusion and jumps of the types illustrated in the previous picture.

The Attack Vector

In the paper we show that the security manager only has one vector stochastic integral to evaluate,

$$X\left(t,T\right)=\int_{t}^{T}a\left(X_{\omega}\left|\theta\right.\right)d\omega \tag{2}$$

We have to now specify a general model that is to be fitted to data Aït-Sahalia et al. (2010) outines a very general model that captures:

- Stochastic volatility in the continuous diffusion.
- Jumps with either deterministic intensity, self exciting intensity and/or self exciting intensity.

The Attack Vector

The attack vector consists of a deterministic drift term $(u_i dt)$, its own volatility term $(V_{i,t})$, and a jump term, dN of size Z.

$$dX_{i,t} = u_i dt + \sqrt{V_{i,t}} dW^X_{i,t} + Z_{i,t} dN_{i,t} \tag{3}$$

where $dW^X_{i,t}$ is a Brownian motion. The volatility equation (4) is given a stationary stochastic process:

$$dV_{i,t} = k_i \left(\theta_i - V_{i,t}\right) dt + \eta_i \sqrt{V_{i,t}} dW^V_{i,t} \tag{4}$$

where $dW^V_{i,t}$ is a Brownian motion, θ_i denotes the long-term volatility, k_i the speed of adjustment, and η_i denotes the kurtosis.

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

$$\begin{cases} \mathbb{P}\left[N_{i,t+\Delta} - N_{i,t} = 0 \,|F_t|\right] = 1 - \lambda_{i,t}\Delta + o\left(\Delta\right) \\ \mathbb{P}\left[N_{i,t+\Delta} - N_{i,t} = 1 \,|F_t|\right] = \lambda_{i,t}\Delta + o\left(\Delta\right) \\ \mathbb{P}\left[N_{i,t+\Delta} - N_{i,t} > 1 \,|F_t|\right] = o\left(\Delta\right) \end{cases} \tag{5}$$

where $N_{i,i+\Delta}$ is an m point process counting the number of jumps in $(0, t+\Delta)$ for the $i = 1, \ldots, m$ 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:

$$\lambda_{i,t} = \lambda_{i,\infty} + \sum_{j=1}^{m} \int_{-\infty}^{t} g_{i,j}\left(t-s\right) dN_{j,s} \tag{6}$$

where $i = 1, \ldots, m$ and $s \leq t$, and $j = 1, \ldots, m$; 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 and $g_{i,j}\left(t-s\right) = \beta_{i,j}e^{-\alpha_i(t-s)}$.

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

$$\begin{aligned}
&\mathbb{E}\left[\Delta X_t\right] = \left(\mu + \lambda M\left[1\right]\right)\Delta + o\left(\Delta^2\right)\\
&\mathbb{E}\left[\left(\Delta X_t - \mathbb{E}\left[\Delta X_t\right]\right)^2\right] = \left(\theta + \lambda M\left[2\right]\right)\Delta + \frac{\beta\lambda\left(2\alpha-\beta\right)}{2\left(\alpha-\beta\right)}M\left[1\right]^2\Delta^2 + o\Delta^2\\
&\mathbb{E}\left[\left(\Delta X_t - \mathbb{E}\left[\Delta X_t\right]\right)^3\right] = \lambda M\left[3\right]\Delta\\
&+\frac{3}{2}\left(\eta\theta\rho^V + \frac{\left(2\alpha-\beta\right)\beta\lambda M\left[1\right]M\left[2\right]}{\left(\alpha-\beta\right)}\right)\Delta^2 + o\left(\Delta^2\right)
\end{aligned} \tag{7}$$

From these moment conditions, plus the Kurtosis and some co-moment conditions we can fit the model to data using the method of moments.

The Data

- For our statistical analysis, we pick DShield data for 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 find missing dates, with missing values filled using piecewise cubic spline interpolation.
- We then compute the individual and multivariate moments for this process and use these estimated moments to derive the parameters for the process with the equivalent moments.

Table: 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 recipientÕs

How to Interpret The Results

- One of the issues with any multivariate model is that the number of parameters explodes and analysis of a ten variate Hawkes process has 130 parameters fitted.
- In keeping with the finance and economics literature in this area we focus on collections of parameters and joint hypotheses.
- Our first set of comparators looks at jumps that have deterministic versus stochastic intensities in jumps.
- Our second analysis asks whether jumps are contagious.
- The final analysis looks at the critical components of the attack vector which excite jumps across the system.

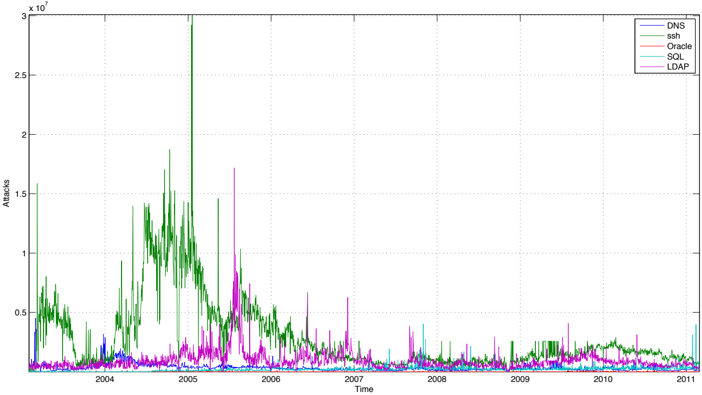
x 10^7
3
2.5
2
1.5
1
0.5
Attacks
2004
2005
2006
2007
2008
2009
2010
2011
Time
DNS
ssh
Oracle
SQL
LDAP

Long Run Intensities

Table: Long-run Intensities; Diagonal Elements of G, for $\tau = 1$ day

	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

Table: Normalized $G(\tau)$ Matrix

	DNS	ssh	Oracle	SQL	LDAP	Web Server	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
Web Server	0.91	0.94	0.76	0.75	0.54	1	0.67	0.84	0.89	0.88
Secure Web Server	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

Conclusions and Future Directions

- To a high level of statistical certainty the attack process is a jump diffusion rather than a simple diffusion with stochastic volatility.
- The jumps almost certainly exhibit stochastic intensities.
- Analysis of the $G(\tau)$ matrix suggests that the jump intensities exhibit both mutual and self excitation properties. Hence contagion across the attack vector.
- The high levels of persistence in this system indicate that jump shocks are most likely permanent over a reasonable time horizon.
- Suggests that a standard mean-variance risk approach for basing cost-benefit calculations is inadequate.

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. Security games with contagion. Manuscript, 2011: `http://www.econ.cam.ac.uk/faculty/goyal/wp11/securitygames17.pdf`.

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.

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*.