Analysis and Modeling for Safer Cyberspace

Research Statement

Benjamin Edwards

November 29, 2016

My research takes the perspective that achieving computer security requires both rigorous empirical measurement and models to understand what defenses and interventions will be most effective. Many cybersecurity problems today occur at a global scale, involving nations, corporations, or individuals whose actions have impact around the world. Despite these global, persistent problems, there is limited research on the actual effectiveness of the many interventions that have been proposed or deployed. For example, botnets have been a vehicle for malicious behavior for more than 15 years, but it is unclear whether the most popular intervention, the botnet takedown, has been effective. Most interventions are inspired by deep, hands-on experience with specific attacks and are never evaluated systematically at a large scale. Moreover, as the scope of cyber-insecurity has increased, no one security practitioner can possibly know all of the relevant details associated with the challenges we face today. Thus, there is a need for more explicit and rigorous methods to determine which interventions are effective and which are are not.

Collection and analysis of large-scale security data is a crucial part of this program. Straightforward analysis will rarely be possible with security relevant datasets. Substantial work will be required to transform unstructured data into meaningful signals. Moreover, relevant measures of security, such as the concentration of infected machines within an organization, is often heavy tailed (varies over many orders of magnitude) making it difficult to separate the effect of interventions from typical fluctuations. I believe that appropriate, rigorous data collection, analysis, and modeling are all needed to secure our increasingly interconnected and computationally reliant society. My research addresses these issues by developing data-driven and abstract models to investigate security phenomena and to test interventions.

1 Cyber Security Informatics

Collecting and analyzing large amounts of data has led to new insights in physics, biology, and economics, but these methods have not been widely applied in security, even as more data on security incidents is becoming available. Data-driven modeling has allowed me to study two widely investigated security phenomena and produce novel results.

1.1 Data Breach Hype and Heavy Tails

In 2015, I examined trends in data breaches over the past decade [2]. Widely publicized data breaches have exposed the personal information of hundreds of millions of people. Some reports point to alarming increases in both the size and frequency of data breaches, spurring institutions around the world to address what appears to be a worsening situation. We studied a popular public dataset of US breaches maintained by Privacy Rights Clearinghouse and developed rigorous statistical models to investigate trends in data breaches. We used Bayesian generalized polynomial trend models to investigate the distribution of data breach size and frequency, and differentiate between different possible trends. Careful statistical analysis showed that neither size nor frequency of data breaches has increased over the past decade. More importantly, we found that the apparent increases that have attracted so much media attention can be explained by the heavy-tailed statistical distributions that best describe the data. We were able to use our model to predict the probability of major breaches in the future, and showed that even without increases in breach frequency or size, the heavy tailed nature of breaches indicates that we can expect more large breaches in the future. We measured the size of breaches based on the number of records they contained, and we extended the model to include findings from other researchers concerning the cost of data breaches. Public data on the true cost of a breach is sparse, and is a promising area for future research. Our work won the best paper award at the Workshop on the Economics of Information Security, and was covered in the security community (Schneier on Security) and the mainstream media (eWeek, Government Technology). Models such as these can help us to understand the nature of security phenomena better before policymakers take steps to address security problems

1.2 Modeling Ten Years of Spam Interventions

I have also used data-driven statistical models to study the impact of popular interventions such as botnet takedowns. Our recent paper at the Annual Computer Security Applications Conference investigated trends in worldwide email spam from a data set of consisting of 127 Billion spam messages sent from 440 million unique IP addresses spread across 260 ISPs in 60 countries over the course of a decade [3]. We showed that geography, national economics, Internet connectivity and traffic flow all impact local spam concentrations. We developed statistically robust time-series models and found similar heavy-tailed, high variance concentrations of spam sending IP addresses. This high variance data can obscure trends and the effects of interventions. We used the model to identify three statistically distinct eras within the ten-year data set. This simple autoregressive model with two transitions provided better predictions than more complex machine-learning techniques such as support vector machines, artificial neural networks, and random forest regression. When the exact date of an intervention is known (as in the case of botnet takedowns), we can use such a model to analyze its impact. We studied twelve different botnet takedowns and found that most had little long-term impact on global spam levels, in some cases even increasing global spam six weeks after the takedown. Moreover, we found that takedowns have highly localized effects. Takedowns that are highly effective in some countries are followed by increases in other countries at a later date. Future work to predict at risk countries ahead of time could improve the efficacy of future takedowns, and indicate which countries would be effective targets for cyber capacity building.

2 Beyond Data-driven Modeling

Despite the growing availability of security data, the data most pertinent to security questions often has not been collected or is privately held. In cases like these it may not be possible to construct data-driven models to study important security questions. In these cases, more abstract models can provide insight.

2.1 Cyber War and Espionage: The Attribution Problem

A persistent question is how strategies for cyber-warfare differ from those in traditional warfare. National security concerns often preclude researchers from accessing data about cyber-warfare and espionage, and conducting experiments on what strategies might best deter future escalation would be irresponsible. In this case, modeling is the only way to provide insight into this new domain of conflict. In collaboration with Robert Axelrod and his student, we have developed a game theoretic model of cyber-conflict. While some lessons from the cold war and traditional conflict apply to the cyber domain, several factors prevent direct applications of these models. Attribution of cyber attacks to state actors is more difficult as digital evidence is often more complex, malleable and prone to manipulation. Moreover, non-state actors such as crime syndicates, terrorist groups, and patriotic hackers can have capabilities comparable to some state actors. Our model indicate that in many scenarios it is rational for countries to tolerate persistent cyber attacks without response. The work is currently under review for publication in the Proceedings of the National Academy of Sciences.

2.2 Making Search Safer

In 2012, we developed a simple Markov model of malware spread through large populations of websites and studied the effect of two interventions that might be deployed by a search provider: blacklisting and lowering the search rankings of infected web pages [5]. The model established the effectiveness of each intervention as well as externalities that might be associated with false positives. Once again, we found that when traffic to different sites is heavy tailed, the effect of interventions can be difficult to identify, and it will be difficult to determine empirically whether certain website interventions are effective. After publication, Google engineers confirmed that sites with a history of infection are ranked lowered, validating our results.

2.3 Regulation and Security Policy

Effective, long-term solutions to security problems will need to be a partnership between users, developers, and policy makers. In addition to my modeling research, I organized and participated in a panel discussing how regulating the software development industry could affect security issues [4].

2.4 Towards a Theory of Security Interventions

Finally, I have begun work on developing a common language and framework for understanding malicious behavior and responses across many complex systems. Though the variety of malicious behaviors is vast, most defensive responses fall into a few generic categories: Observe, Hide, Filter, Repair and Counterattack. We have successfully identified examples of these types of defenses across many biological, social and technical systems. A common language and framework will support the development and communication of generic models for these processes [1].

3 Future Work

3.1 Cyber-capacity Building for at risk Nations

My work on studying interventions which target spam indicated that botnet takedowns can result in localized increases in spam concentrations. In the future, I plan to work to identify at risk countries and ISPs. Once these at risk entities can be identified, I can investigate which cyber capacity building interventions can help to reduce negative externalities.

3.2 Data Breach Risk and Costs

My work has outlined how to identify trends in data breaches. Two important next steps are to identify what features of an organization might make it susceptible to future breaches, and estimating the cost of breaches. Factors such as internal security policy, employee training, and types of information held may contribute to an organizations risk of breach. The cost of a data breach can be highly variable, and likely depends on more than just the total number of records including: the type of information exposed, what purpose the information was used for, the time taken to discover the breach, and the cost of correcting the security problems that lead to the breach. I have been working with industry partners to obtain data which will be used to build rigorous models of data breaches and their costs.

3.3 Diversity and Insecurity

Using biological analogies, security researchers have argued that a software monoculture poses a risk to users, and a more diverse software ecosystem would prevent widespread vulnerabilities. I am currently working with industry partners to investigate whether software diversity affects the security profile of an organization.

References

[1] Benjamin Edwards and Stephanie Forrest. How do complex systems protect themselves form malicious behavior. *Conference on Complex Systems*, September 2015.

[2] Benjamin Edwards, Steven Hofmeyr, and Stephanie Forrest. Hype and heavy tails: A closer look at data breaches. *Workshop on the Economics of Information Security 2015*, June 2015.

[3] Benjamin Edwards, Steven Hofmeyr, Stephanie Forrest, and Michel van Eeten. Analyzing and modeling longitudinal security data: Promise and pitfalls. *Annual Computer Security Applications Conference*, December 2015.

[4] Benjamin Edwards, Michael Locasto, and Jeremy Epstein. Panel summary: The future of software regulation. In *Proceedings of the 2014 workshop on New Security Paradigms Workshop*, pages 117–126. ACM, 2014.

[5] Benjamin Edwards, Tyler Moore, George Stelle, Steven Hofmeyr, and Stephanie Forrest. Beyond the blacklist: modeling malware spread and the effect of interventions. In *Proceedings of the 2012 workshop on New security paradigms*, pages 53–66. ACM, 2012.