Uncertainty in Interdependent Security Games*

Benjamin Johnson[a], Jens Grossklags[b], Nicolas Christin[a], and John Chuang[c]

[a]CyLab, Carnegie Mellon University
[b]Center for Information Technology Policy, Princeton University
[c]School of Information, UC Berkeley
{johnsonb,nicolasc}@andrew.cmu.edu
jensg@princeton.edu
chuang@ischool.berkeley.edu

Abstract. Even the most well-motivated models of information security have application limitations due to the inherent uncertainties involving risk. This paper exemplifies a formal mechanism for resolving this kind of uncertainty in interdependent security (IDS) scenarios. We focus on a single IDS model involving a computer network, and adapt the model to capture a notion that players have only a very rough idea of security threats and underlying structural ramifications. We formally resolve uncertainty by means of a probability distribution on risk parameters that is common knowledge to all players. To illustrate how this approach might yield fruitful applications, we postulate a well-motivated distribution, compute Bayesian Nash equilibria and tipping conditions for the derived model, and compare these with the analogous conditions for the original IDS model.

1 Introduction

Starting with the Morris Worm in 1988, security attacks on computer systems have gradually shifted from "point-to-point" attacks, where a single attacker targets a single defender, e.g., to deny service, to propagation attacks, where the attacker attempts to compromise a few machines and, similar to an epidemic, uses these compromised machines ("bots" or "zombies") to infect additional hosts. The advantage of propagation attacks is that the miscreants behind them can commandeer reasonably quickly a very large pool of machines, which can, in turn, be monetized. Among many other activities, bots have been used to send spam email, host phishing websites [22], or acquire banking credentials [5].

Traditional security models that pit a defender (or a set of defenders) against an external attacker may not capture all the intricacies of propagation attacks, as the attacker population may vary over time. In contrast, models of interdependent security

* This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by the National Science Foundation through award CCF-0424422 (TRUST - Team for Research in Ubiquitous Secure Technology).

(e.g., [18]), where hosts in the network may (involuntarily or not) act on behalf of the attacker, appear more suitable to characterize propagation attacks.

Interdependent security models have been used in the context of airline security [14], and disease propagation [15]. In these contexts, it may be possible to characterize infection rates or measure attack probability based on historical data. In the context of information security, on the other hand, we posit that uncertainty on the possibility of an attack, and ambiguity on the configuration of other networked hosts imposes significant challenges for the selection of effective security strategies.

For instance, networks in many organizations may be quite large, and are prone to have poorly known configuration parameters, even by their own administrators [19]. A firewall that governs the entrance to the network may have thousands of rules, some of them obsolete, some of them redundant, and thus it may be difficult to explicitly characterize the probability a given outside attack could actually succeed in penetrating the corporate network. Network configurations may be relatively complex, and two machines located close to each other geographically may be far apart in the network topology. In the end, network administrators may only have very rough estimates of the various probabilities of external attacks or of attack propagation between interior nodes [3].

The contribution of this paper is to introduce and exemplify a method for resolving risk uncertainty by means of a well-motivated probability distribution on risk parameters. We introduce the method within the context of a single interdependent security game that draws its motivation from an organizational LAN in which agents have a significant residual impact on the security of their own and their peers' resources (e.g., such as in university and many corporate networks). Our examples show that such distributions can be easily motivated, and that the resulting derived conditions for equilibria and tipping effects are reasonable, in the sense that they compare similarly to equilibrium conditions derived in the original IDS model using the distribution's expected values of the model's risk parameters.

The rest of this paper is organized as follows. We review related work in Section 2. In Section 3, we describe our model, which is directly inspired by the work of Kunreuther and Heal [18], and explain how we address risk uncertainty within this model. We provide formal and numerical analysis of interdependent security games with homogeneous and heterogeneous populations, and with or without uncertainty, in Section 4. We conclude in Section 5.

2 Related work

2.1 Interdependent security

In their 2003 study, Kunreuther and Heal formalize the concept of interdependent security with their primary example stemming from the airline industry [14,18]. In this case, the individual airlines are concerned about a major single attack that may originate at

some point in the network, but could be propagated to another airline in the system. Airlines can defend themselves against direct attacks, however, they are powerless against dangerous loads received from other aviation entities. In follow-up work, they also consider a game in which players can protect themselves effectively against direct and indirect attacks through some protection measure (e.g., vaccination), however the benefit of the security investment diminishes with its popularity in the population [15]. This research has motivated follow-up contributions in algorithmic computation of equilibria with real-world data [17], and human-subjects experimentation in the laboratory [16].

In this paper, we refer to a complementary computer security model commented upon by Kunreuther and Heal [18]. In this scenario, a single compromised network resource can adversely impact other connected entities multiple times. We study this game more formally by deriving game-theoretic equilibrium solutions for different information conditions and network-wide behaviors (e.g., tipping point phenomena).

Concurrently to the research on interdependent security, Varian started a formal discussion on the role of security as a public good [25]. In our work, we expanded on his work by developing a security games framework including additional games and investment strategies (i.e., self-insurance) [9]. We also considered the impact of player heterogeneity [10], and the influence of strategically acting attackers on the security outcome [7].

An alternative optimization approach is pursued by Miura-Ko *et al.* who derive Nash equilibrium conditions for simultaneous move games in which the heterogeneous interactions of players can be represented with a set of piece-wise linear conditions [21]. They further enrich their basic model to develop three studies on password security, identity theft, and routing path verification. The authors verify the robustness of their approach to perturbations in the data, however, do not formally consider the role of uncertainty.

2.2 Uncertainty and security

In the context of the value of security information, research has been mostly concerned with incentives for sharing and disclosure. Several models investigate under which conditions organizations are willing to contribute to an information pool about security breaches and investments when competitive effects may result from this cooperation [8]. Empirical papers explore the impact of mandated disclosures [4] or publication of software vulnerabilities [24] on the financial market value of corporations.

Other contributions to the security field include the computation of Bayesian Nash outcomes for an intrusion detection game [20], security patrol versus robber avoidance scenarios [23], and the preservation of location privacy in mobile networks [6]. A different approach is followed by Alpcan and Başar who present an application of game theory and stochastic-dynamic optimization to attack scenarios in the sensor network context [2].

In our prior work, we studied the impact of uncertainty in three different games [11,13]. We also developed a set of metrics to study the value of better information [12].

A more extended review of theoretical and empirical work is presented by Acquisti and Grossklags in which they discuss the moderating role of risk, uncertainty and ambiguity in the areas of privacy and security [1].

3 Model

3.1 Interdependent network security

We focus our attention on interdependent security games that directly model network security. For the basic setup, suppose that each of n players is responsible for operating her personal computer, and that players' computers are connected to each other through a given internal network, e.g., a corporate LAN. Each computer is also connected to an external network, e.g., the Internet. The external connection poses certain risks (e.g., infection with viruses), and if a user's resources are compromised then she will suffer a total loss, normalized to 1. In addition, some of these viruses have the ability to propagate through the internal network to compromise all the other players' computers. If a player's computer is compromised in this way, she also faces a total loss.

Each player has a choice of investing in security mechanisms with a cost c to eliminate the risk of being infected by an external virus. However, there is no effective way to protect from the risk of an internal contamination as the result of another player passing along a virus through the internal network. This modeling choice reflects a relatively common situation in corporate networks where security policies are set to have computers almost blindly trust contents coming from inside the corporate network (which facilitates automated patching, and software updates for instance), while contents coming from outside of the network are thoroughly inspected.

In addressing the risk factors associated with the virus infection and contamination, we consider two versions of this game – a homogeneous version and a heterogeneous version. In the homogeneous version, p is the probability that a given computer becomes infected with a virus, and q is the probability that a computer with a virus contaminates other computers in the system. Since a computer can only transmit a virus once it is infected, we may assume that $q \leq p$. In the heterogeneous version, p_{ii} is the probability that player i becomes infected with a virus, and p_{ij} is the probability that player i causes player j to become contaminated due to virus transmission. Again since a computer must be infected before contaminating another computer, we may assume that $p_{ij} \leq p_{ii}$ for every i and j.

The utility of each player in this game depends not only on her choice to protect, but also on the choices of other players. If there are k players in the network who are not protecting, then player i's choice can be framed as follows. If she protects, then she pays a cost c, eliminating the risk of a direct virus infection, but she still faces the risk of internal contamination from k different players. If she fails to protect, then she does

not pay c, but she faces both the risk of an internal contamination from one of the k players, as well as the risk of an external infection. The utility function for player i is derived directly from these considerations.

For the homogeneous version of the game, the expected utility of player i is given by the equation:

$$U_i = \begin{cases} -c + (1-q)^k & \text{if player } i \text{ protects} \\ (1-p)(1-q)^k & \text{if player } i \text{ does not protect} \end{cases} \tag{1}$$

where k is the number of players other than i who choose not to protect.

In the heterogeneous version of the game, the expected utility for player i is given by:

$$U_i = \begin{cases} -c + \prod_{j \neq i : e_j = 0} (1 - p_{ji}) & \text{if player } i \text{ protects} \\ (1 - p_{ii}) \prod_{j \neq i : e_j = 0} (1 - p_{ji}) & \text{if player } i \text{ does not protect} \end{cases} \tag{2}$$

where e_j in a binary indicator variable telling us whether player j chooses to protect.

3.2 Uncertainty

In the usual treatment of interdependent security (IDS) games such as the one above, the risk parameters are assumed to be known. We are interested in the case in which the risks of virus infection and contamination are unknown. Such uncertainty is especially well-motivated in the IDS computer network game since computer users in general do not know or understand well the potential risks posed by various types of viruses.

For our model with uncertainty, we assume that players do not know the risks, but they believe and agree upon some probability distribution over risk parameters. In other words, there is a probability distribution D that describes players beliefs about the relevant risks. True to rational Bayesian form, everyone believes that the relevant risk parameters are drawn from the same distribution D.

In the homogeneous case, D is a distribution on $[0,1] \times [0,1]$, representing the players' mutually-held beliefs about the parameters p and q. In the heterogeneous case, D is a distribution on $[0,1]^{n \times n}$, representing players' mutually-held beliefs about the parameters p_{ij}.

4 Analysis

4.1 Overview

Our analysis focuses on determining equilibrium conditions. We start with the homogeneous version and then proceed to the heterogeneous version. In each case, we begin by looking at the game with full information and computing conditions under which various Nash equilibria exist and how they can be tipped or disrupted. We extend these

results to the realm of uncertainty by positing a general distribution D and rewriting the equilibria conditions using expected values of aggregate risk parameters conditioned on D. We follow by providing and motivating a parametrized example distribution D_ϵ and using this distribution to compute various equilibrium conditions explicitly. In the homogeneous version, we analyze these conditions numerically and graphically, and compare the results to the original IDS game in which risk parameters are known.

4.2 Homogeneous case: A monoculture of potential failure modes

Nash Equilibrium We begin with the homogeneous case. Let's first assume that p and q are known. This game has two possible strong Nash equilibria, one in which all players protect, and one in which no player protects. Considering a simple cost-benefit analysis, the "everyone-protects" equilibrium is achievable if and only if the cost of protection is less than the cost of an external infection (i.e. $c < p$). Similarly, the "everyone-defects" equilibrium is achievable if and only if the cost of protection is greater than the likelihood that a player is infected, but not compromised, assuming that all of the other players are failing to protect (i.e. $c > p(1-q)^{n-1}$). In the middle area $p(1-q)^{n-1} < c < p$, both equilibria are possible, the protection equilibrium is Pareto optimal, and both equilibria are subject to the possibility of tipping phenomena in which forcing a certain number of players to switch strategies will effect the opposite equilibrium.

Tipping Phenomenon To understand this game's tipping phenomenon when there are n players, it suffices to understand the game's defection equilibrium conditions when there are k players and $k < n$.

If players are in an "everyone defects" equilibrium, then to tip the equilibrium to one in which everyone protects, it is necessary (and sufficient) to force protection upon enough players so that universal defection among the remaining players is no longer an equilibrium strategy. The number of forced protections required to accomplish this is the least integer k such that $c < p(1-q)^{n-1-k}$. In words, k is the least integer such that universal defection fails to be an equilibrium strategy in a game with only $n-k$ players.

Similarly, if players are in an "everyone protects" equilibrium, then the number of defections required to tip the equilibrium toward universal defection is the least integer k such that $c > p(1-q)^k$. In this case k is the least integer such that, in a game with $k+1$ players, universal defection is an equilibrium strategy.

In any case, the boundary conditions that describe the tipping phenomenon are the same conditions that describe defection equilibria in games with fewer players.

Uncertainty When dealing with a joint probability distribution over the parameters p, q, the above reasoning applies with the exception that players compute an expected

value for p and $p(1-q)^{n-1}$ using the distribution D. Thus "everyone protects" is an equilibrium if and only if $c < E_D[p]$ and "everyone defects" is an equilibrium if and only if $c > E_D[p(1-q)^{n-1}]$. The tipping phenomenon have an analogous translation involving these expected values.

Example distribution D_ε To exemplify the scenario, we propose a class of distributions D_ε, parametrized by a number $\varepsilon \in [0,1]$. To motivate this distribution, we suppose that there is a fixed number $\varepsilon \in [0,1]$ such that players believe the risk of external infection, p, is no more than ε. D_ε then assigns a probability to the pair $p, q \in (0,1) \times (0,1)$ according to the following two-step procedure. First draw p from the uniform distribution on $(0, \varepsilon)$. Then draw q from the uniform distribution on $(0, p)$. Since the only thing players really know for certain about the risks are that $0 \leq q \leq p \leq 1$, the parametrized distribution D_ε represents an effort to reflect the notion that "infection is somewhat unlikely, ('somewhat' being explicitly quantified by the parameter ε), and contamination as a result of infection is even less likely, and aside from that we do not have a very good idea what the risk is."

Bayesian Nash equilibrium for D_ε To determine the Bayesian Nash equilibrium conditions for the parametrized game with uncertainty, we must compute the expected values $E_{D_\varepsilon}[p]$ and $E_{D_\varepsilon}[p(1-q)^{n-1}]$ explicitly. The expected value of p under D_ε is $\frac{\varepsilon}{2}$, because p is drawn from the uniform distribution on $(0, \varepsilon)$. The expected value of $p(1-q)^{n-1}$ under D_ε can be computed by evaluating the expression:

$$\frac{1}{\varepsilon} \int_0^\varepsilon \left(\frac{1}{p} \int_0^p p(1-q)^{n-1} dq \right) dp \tag{3}$$

where the inner integral is to be evaluated assuming that p is constant relative to q. The expression evaluates to

$$\frac{1}{n} \left(1 - \frac{1-(1-\varepsilon)^{n+1}}{\varepsilon(n+1)} \right). \tag{4}$$

When $\varepsilon = 1$ this expression simplifies to $\frac{1}{n+1}$. In practical terms, the parameter selection $\varepsilon = 1$ describes a situation in which players have so little knowledge of the risk factors, that they may as well believe the parameters are uniformly distributed across all possible options. Under such conditions and with many players, the protection costs must be very small to counteract defection incentives. On the other hand, from a social planner's point of view the situation may be manageable, as the total cost (cost per player $\times$ number of players) necessary to properly incentivize network protection is bounded by a constant independent of the network size.

Graphical analysis Figure 1 plots the boundary conditions for Bayesian Nash equilibrium as a function of ε, for a range of N. For comparison, Figure 2 plots the boundary

conditions for Nash equilibria in the full information case when $p = \frac{\varepsilon}{2}$ and $q = \frac{\varepsilon}{4}$. Figure 3 exemplifies the equilibrium tipping phenomenon in a 7-player game.

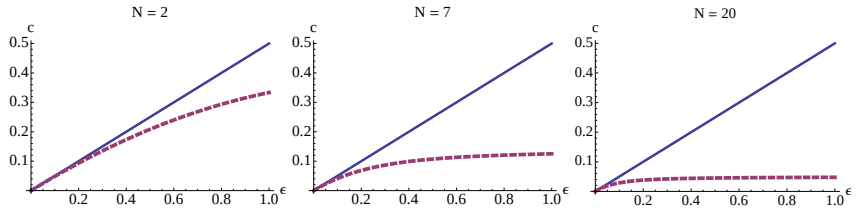

Fig. 1. Bayesian Nash equilibrium boundaries for the homogeneous game with N players. If (c, ε) is below the solid line then "everyone protects" is a Bayesian Nash Equilibrium. If (c, ε) is above the dashed line, then "everyone defects" is a Bayesian Nash equilibrium. In the middle area, there are competing equilibria, and tipping points.

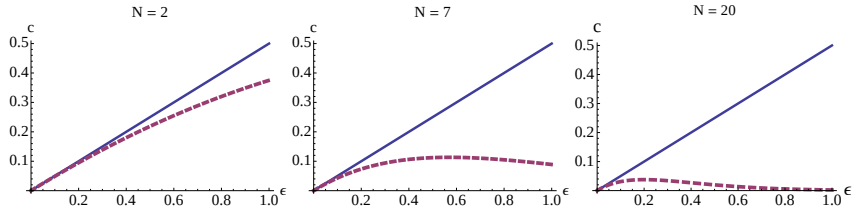

Fig. 2. Nash equilibrium boundaries for the homogeneous case with p and q common knowledge among all players. For comparison with Figure 1, we assume that $p = \frac{\varepsilon}{2}$ and $q = \frac{\varepsilon}{4}$.

4.3 Heterogeneous case: Unknown and diverse configuration problems

Nash Equilibrium For the heterogeneous case, we begin by assuming the p_{ij} are known. Here, once again, the strategy "everyone protects" is a Nash equilibrium if and only if $c < p_{ii}$. The strategy "everyone defects" is a Nash equilibrium if and only if $c > p_{ii} \prod_{j \neq i}(1 - p_{ji})$. In the middle area $p_{ii} \prod_{j \neq i}(1 - p_{ji}) < c < p_{ii}$, both equilibria are possible, the protection equilibrium is Pareto optimal, and the situation is subject to the tipping phenomenon.

Tipping Phenomenon The tipping phenomenon in the heterogeneous case is completely analogous to the homogeneous case. Tipping conditions for an n-player game

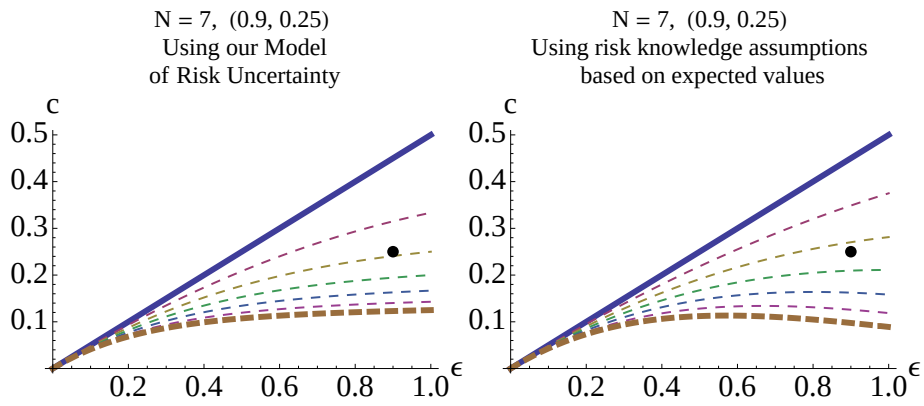

Fig. 3. Tipping point boundaries for the homogeneous game with 7 players. In this example, the risk threshold value is $\varepsilon = 0.9$ and the protection cost is $c = 0.25$. In the model incorporating uncertainty, it takes 2 defections to tip a protection equilibrium into a defection equilibrium, while in the model in which risk parameters are known, it takes 3 defections to tip the equilibrium from protection to defection. In the other direction from universal defection to universal protection, tipping the equilibrium requires 5 forced protections in the case of the uncertainty model and 4 forced protections in the case of knowledge assumptions.

are determined by considering the defection equilibrium conditions for games with fewer players.

Uncertainty In the presence of uncertainty with beliefs about risk parameters governed by a general distribution D, the same analysis as above holds with $E_D[p_{ii}]$ and $E_D[p_{ii}\prod_{i\neq j} p_{ji}]$ replacing p_{ii} and $p_{ii}\prod_{j\neq i}(1-p_{ji})$, respectively.

Example distribution D_ε To exemplify the heterogeneous scenario, we propose a class of distributions D_ε, analogous to the homogeneous case, again parametrized by a number $\varepsilon \in [0,1]$. As before, players believe the risk of external infection, p_{ii}, is no more than ε. D_ε then assigns a probability to the matrix $p_{ij} \in (0,1)^{n\times n}$ according to the following procedure. First draw each p_{ii} independently from the uniform distribution on $(0,\varepsilon)$. Then draw each p_{ij} independently from the uniform distribution on $(0,p_{ii})$, so that $0 \leq p_{ij} \leq p_{ii} \leq 1$ for every $i \neq j$.

Bayesian Nash equilibrium for D_ε To determine the Bayesian Nash equilibrium conditions for the parametrized game with uncertainty in the heterogeneous case, we must compute the expected values $E_{D_\varepsilon}[p_{ii}]$ and $E_{D_\varepsilon}[p_{ii}(1-p_{ji})^{n-1}]$ explicitly. Unlike the homogeneous case, these expected values are trivial to compute because all the vari-

ables in relevant expressions are independent, thus we can use linearity of expectation. The expected value of p_{ii} is $\frac{\varepsilon}{2}$, and the expected value of $p_{ii}(1-p_{ji})^{n-1}$ is $\frac{\varepsilon}{2}(1-\frac{\varepsilon}{4})^{n-1}$.

We omit the graphical analysis for the heterogeneous case both due to space constraints and because there is no simple way to compare results with the original model due to differences in the number of free parameters.

Another example distribution, $D_{\epsilon,i}$ One final example to consider is one in which players mutually acknowledge that some computers are more likely to be infected than others. We can exemplify this scenario by using a distribution $D_{\varepsilon,i}$ that discriminates among risk parameters for different players. $D_{\varepsilon,i}$ assigns a probability to the matrix $p_{ij} \in (0,1)^{n\times n}$ according to the following procedure. First draw each p_{ii} independently from the uniform distribution on $(0,\varepsilon_i)$. Then draw each p_{ij} from the uniform distribution on $(0,p_{ii})$. The distribution $D_{\varepsilon,i}$ reflects the same uncertain sentiment regarding risk as D_{ε}, yet it also accommodates a notion – certainly realized in practice – that some assets bear higher risk level than others.

Under the distribution $D_{\varepsilon,i}$, the computations involved in determining each player's strategic response to the behavior of others are analogous to those computations under the distribution D_{ε}. Again the individual variables in the relevant expressions are drawn independently so that linearity of expectation can be applied. For example, when all other players are failing to protect, player i will also fail to protect if and only if $c > \frac{\varepsilon_i}{2}\prod_{j\neq i}(1-\frac{\varepsilon_j}{4})$. Unfortunately, determining all possible Bayesian Nash equilibria requires addressing a number of caveats, because players have different incentives due to the homogeneity in their beliefs about their respective risks. We defer a thorough analysis of this scenario to future work.

5 Discussion and Conclusions

Interdependent models of information security in corporate networks seem especially well-motivated, but it is difficult to utilize the sharpness of these models due to uncertainty regarding real world risk factors. Our approach has been to make these models smoother, by incorporating players' uncertainty about various risk parameters.

Our objective has been to develop a mechanism for dealing with risk uncertainty in a security context. We focused on a single IDS model involving a computer network, and we adapted the model to capture a notion that players have only a very rough idea of security threats and underlying structural ramifications. We formally resolved this uncertainty by means of a probability distribution on risk parameters, one that was common knowledge to all players. We postulated a reasonable such distribution, computed Bayesian Nash equilibria and tipping conditions for the resulting model, and compared those to the same conditions for the original model.

Crucially from a practical standpoint, we incorporated this new probabilistic machinery while actually assuming less – indeed our adapted model using the example

distribution D_ε reduced the number of free parameters. Nonetheless, we found that the adapted model maintains characteristic equilibrium properties and asymptotic behaviors when information assumptions are relaxed. There are still only the two extreme equilibria. There is still a range of cost and risk distribution parameters for which the equilibrium can be tipped the other way by encouraging some players to switch strategies. Even the boundary conditions for equilibrium conditions and tipping effects are similar to those obtained from the original model, and we would expect such similarities to extend to other well-motivated probability distributions in other contexts.

There were some mild differences compared to the full knowledge model using the distribution's expected values of model parameters. In our homogeneous model incorporating uncertainty, a generally low contamination risk facilitated the possibility of slightly more defections, while a generally moderate to high contamination risk facilitated fewer defections. An application of this phenomenon is that when risks are small, it may be better from a social planner's standpoint to communicate such risks by using expected values of parameters, while if risks are large it may be better to present them in a manner that incorporates uncertainty using a distribution.

As a general rule, when we apply a security model to a real world situation, we expect that some real world data will be substituted for the parameters in the model. Unfortunately this is oftentimes difficult or impossible to do, especially for risk parameters. Without knowing the risks, we are left with the problem of how to use the model for anything at all. Our approach addresses this situation in a reasonable way for a very simple model. The approach itself is quite general and we expect to find additional applications in future work.

References

1. A. Acquisti and J. Grossklags. What can behavioral economics teach us about privacy? In A. Acquisti, S. Gritzalis, S. Di Vimercati, and C. Lambrinoudakis, editors, *Digital Privacy: Theory, Technologies, and Practices*, pages 363–380. Auerbach Publications, Boca Raton, FL, 2007.
2. T. Alpcan and T. Basar. An intrusion detection game with limited observations. In *Proceedings of the 12th International Symposium on Dynamic Games and Applications*, Sophia Antipolis, France, July 2006.
3. M. Bashir and N. Christin. Three case studies in quantitative information risk analysis. In *Proceedings of the CERT/SEI Making the Business Case for Software Assurance Workshop*, pages 77–86, Pittsburgh, PA, September 2008.
4. K. Campbell, L. Gordon, M. Loeb, and L. Zhou. The economic cost of publicly announced information security breaches: Empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448, 2003.
5. J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of 14th ACM Conference on Computer and Communications Security (CCS)*, pages 375–388, Alexandria, VA, October 2007.

6. J. Freudiger, M. Manshaei, J.-P. Hubaux, and D. Parkes. On non-cooperative location privacy: A game-theoretic analysis. In *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS'09)*, pages 324–337, Chicago, IL, November 2009.
7. N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In *Proceedings of the 2009 Financial Cryptography Conference (FC'09)*, pages 167–183, Accra Beach, Barbados, January 2009.
8. E. Gal-Or and A. Ghose. The economic incentives for sharing security information. *Information Systems Research*, 16(2):186–208, June 2005.
9. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 2008 World Wide Web Conference (WWW'08)*, pages 209–218, Beijing, China, April 2008.
10. J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*, pages 160–169, Chicago, IL, July 2008.
11. J. Grossklags and B. Johnson. Uncertainty in the weakest-link security game. In *Proceedings of the International Conference on Game Theory for Networks (GameNets 2009)*, pages 673–682, Istanbul, Turkey, May 2009.
12. J. Grossklags, B. Johnson, and N. Christin. The price of uncertainty in security games. In *Proceedings (online) of the Eighth Workshop on the Economics of Information Security (WEIS)*, London, UK, June 2009.
13. J. Grossklags, B. Johnson, and N. Christin. When information improves information security. In *Proceedings of the 2010 Financial Cryptography Conference (FC'10)*, pages 416–423, Tenerife, Spain, January 2010.
14. G. Heal and H. Kunreuther. IDS models of airline security. *Journal of Conflict Resolution*, 49(2):201–217, April 2005.
15. G. Heal and H. Kunreuther. The vaccination game. Technical report, Columbia Business School & The Wharton School, January 2005.
16. R. Hess, C. Holt, and A. Smith. Coordination of strategic responses to security threats: Laboratory evidence. *Experimental Economics*, 10(3):235–250, September 2007.
17. M. Kearns and L. Ortiz. Algorithms for interdependent security games. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 561–568. MIT Press, Cambridge, MA, 2004.
18. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, March 2003.
19. F. Le, S. Lee, T. Wong, H. Kim, and D. Newcomb. Detecting network-wide and router-specific misconfigurations through data mining. *IEEE/ACM Trans. Netw.*, 17(1):66–79, 2009.
20. Y. Liu, C. Comaniciu, and H. Man. A Bayesian game approach for intrusion detection in wireless ad hoc networks. In *Proceedings of the Workshop on Game Theory for Communications and Networks*, Pisa, Italy, October 2006.
21. A. Miura-Ko, B. Yolken, J. Mitchell, and N. Bambos. Security decision-making among interdependent organizations. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 66–80, Pittsburgh, PA, June 2008.
22. T. Moore, R. Clayton, and R. Anderson. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, Summer 2009.

23. P. Paruchuri, J. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games. In *Proceedings of the 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 895–902, Estoril, Portugal, May 2008.
24. R. Telang and S. Wattal. An empirical analysis of the impact of software vulnerability announcements on firm stock price. *IEEE Transactions on Software Engineering*, 33(8):544–557, 2007.
25. H.R. Varian. System reliability and free riding. In L.J. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.