Cybersecurity Investment Guidance:
A Note on Extensions of the Gordon and Leob model

Scott Farrow and Jules Szanton[1]
UMBC/CREATE, UM Carey Law/UM CHHS

Abstract

Extensions of the Gordon-Leob (GL) model are first based on mathematical equivalency with a generalized homeland security model. Extensions include limitations on changes in the probability of attack, simultaneous effects on probability and loss, diversion of attack, and shared non-information defenses. The importance of negative external effects which increases the optimal cyberinvestment is investigated through case studies in legal liability. The generality of limitations on cyberinvestment appear limited but require empirical work for calibration.

1. Introduction

The most pressing cyberthreats once came from emailed viruses, today's cyberattacks increasingly take the form of massive identity and intellectual property thefts and a potential for physical damage to critical infrastructure. As cyberattacks have proven to

[1] Appreciation is extended to Larry Gordon, Anupam Joshi, Marty Leob Markus Rauschecker and Matt Shabat for comments and to Andrew Naviasky for research assistance; to the University of Maryland Center on Health and Homeland Security for coordination of the legal research and the Department of Homeland Security Cyber Security and Communications division for short term space. Funding was provided by the Department of Homeland Security National Center on the Risk and Economics of Terrorism Events.

be increasingly disruptive to the economy, a growing body of scholarship examines how much should firms should invest in protection and what are appropriate roles for governments. Gordon and Leob (GL, 2002, 2006) and later, Gordon, Leob, Lucyshyn and Zhou (GLLZ, 2015) have been leaders in examining the optimal level of spending that organizations should invest in cybersecurity. Their optimization approach is an unconstrained expected profit maximization model where cybersecurity investments are separable from other activities of the firm. The "profit" or firm benefit from investing in cybersecurity is thus a cost reduction; the remaining probability of a security breach[2] (S(z)) times the loss (l), which can be altered based on investing in cybersecurity, z. GL discussed damages as direct (private), while GLLZ extended the notation such that social damages are the sum of private damages and external damages consistent with common notation in economics.

GL and GLLZ investigate the implications of their model in some detail after first deriving the expected condition that the optimum (interior) investment is found by determining where the incremental benefits of information security equal the incremental costs. As the optimal investment is shown to be increasing in damages (losses), including external damages increases the optimal level of investment as expected. As with standard models of investment, an organization that only considers private losses in its optimization is correct if there are no external losses; but if external losses exist then the optimum social expenditures increase. By investigating several functional forms for the security breach function, GL and later researchers showed that it is not uncommon for investments to have a maximum of about 37 percent expected losses although this result is specific to the specification.

[2] The "probability" appears to be a truncated density function with maximum v and minimum zero.

An important part of the homeland security literature focuses on security in response to an intelligent adversary. Such a framing leads quickly to game theoretic modeling where each player's move takes into account the actions of other players. The role of state actors in cybersecurity may suggest the potential for such modeling in some circumstances. The approach taken here is that some concerns of game theory are incorporated in three ways: 1) the initial probability of a breach is not constant across sites but reflects the interests of the attacker, 2) the change, if any, in probability of a security breach empirically reflects the (possibly game theoretic) behavioral response of the attacker, and 3) limitations on the ability to reduce the probability of a breach may reflect the characteristics of the attacker the system being protected. For instance, Hausken (2014) provides a cybersecurity example how low level defenses may deter low level attackers and imply a convex security breach function; while also investigating more skilled attackers and more valuable information sets which can imply a concave breach function with different investment implications.

2. *The Equivalency of GL and a General Homeland Security Investment Model*

While GL examine how much an individual firm should invest in preventing a cyberattack, a related body of scholarship examines how much the government should invest in homeland security. In the years following the September 11, 2001 terrorist attacks, set of foundational expected, constrained cost minimization models for homeland security expenditures was published following internal Government development (Farrow, 2007). The GL and Farrow general models make similar assumptions about continuity and derivatives of key functions[3]. The core similarities and differences are investigated below including a number of extensions presented in Farrow. Hence this note focuses on general investment rules where there are different types of constraints

[3] Hausken (2014) discusses how the skills of the attackers may affect the convexity (leading to an interior solution), or the concavity (leading to a corner solution) of the problem.

and issues such as multiple sites, budget constraints, the diversion of attack probability, a minimum attack probability as from Advanced Persistent Threats, public or shared defensive methods, damage reducing investments and other variations.

The core model definitions for Farrow are as below. The “organization” was originally modeled as the government, since governments are expected to consider external effects and select the socially optimal outcome. The organization could also be a firm or consumer. It must be noted, of course, that the government has different incentives than an individual firm or consumer. While the government is concerned with minimizing overall social costs, a self-interested firm or consumer will not consider the external effects of its choices unless there is some feedback mechanism, such as legal liability, which incentives the firm or individual to internalize the costs it creates for others.

Define:

$e_{i:}$: organizational expenditures on site i
$\overline{E}$: a chosen aggregate expenditure level over all sites and pathways.
$P(e_{i:})$: probability of an event. P’<0; P’’>0 where P’ is the partial derivative and functions are assumed to be twice continuously differentiable. This assumes some behavior or reaction function on the part of the attacker such that expenditures could alter their choice of targets or the expenditures could lead to capture prior to an attack.
$S(e_i)$: non-organizational costs of the investment expenditures such as time or changes in productivity, expect S’>0.
$C(e_{i:})$: social cost given an event happens, C’<0; C”>0, which includes direct costs to the organization, $C^D(e_{i:})$ and costs external to the organization (external costs) $C^E(e_{i:}$. Note that the expenditure amount e_i is always assumed to be obligated and spent whether or not an attack occurs. It is the cost, C, that is conditional on the event occurring.

The organization’s investment problem is stated as choosing the level of expenditure at each site ($e_i^* \geq 0$) in order to minimize expected cost:

Min $$\sum_{i=1}^{N} P(e_i)[C(e_i)+e_i+S(e_i)]+[1-P(e_i)](e_i+S(e_i))$$

Subject to: $$\sum_{i=1}^{N} e_i = \overline{E}$$

$$e_i \geq 0$$

The unconstrained minimization form of the problem is that used by Baryshnikov (2012) in his extension of the underlying mathematical properties of the GL model and by Gordon, Leob and Lucyshyn (2003). GL do not consider the added complexity of a budget constraint, although they note that conflict between the Chief Information Officer and the Chief Executive Officer may affect the derivation of the optimal amount. In the budget constrained problem, a budget larger than the optimum expenditure yields the unconstrained solution (the constraint is not binding), while if there is a budget constraint where security expenditures are less than the unconstrained optimum the Farrow results follow explicitly. Farrow focused on the organization being the government who was assumed to consider both external social costs of an investment and external costs of a terrorist event, C. Numerous variations were investigated for changes in the necessary conditions for optimization. Each model type has information to inform the other.

The GL model can be seen as the dual of the Farrow model, with additional constraints (Baryshnikiov, 2012; Gordon, Leob and Lucyshyn, 2003). Maximizing cost reductions (profit) is the dual of minimizing costs for an interior solution. Relaxing those constraints as was done in Farrow leads to additional insights for cybersecurity. At the same time, the parameterization of GL adds greater interpretation to the Farrow results.

The notation changes and constraints to place the GL model in the Farrow notation are as below.

Table 1: Notation changes and equivalencies

Underlying concept	GL original notation	Farrow notation	Note
Sites or information set; unit of analysis	One unit	Multiple sites i	i=1 for equivalence
Investment (expenditure) in cybersecurity	Z	e	Equivalent
Probability of successful attack given an investment level	S(z) Security breach function	P(e) Terrorism event function	Equivalent (but modified initial limits of prob. Not incorporated in Farrow)

Conditional Cost or Loss of an event	L
	Initially, L; later L^P+L^E for private and external
		C(e)
		Includes both direct and external, later C^P and C^E
		C= L^P+L^E
			GL is constant
			GL start with L^P, Farrow starts with C= C^P+C^E= L^P+L^E and as influenced by expenditures

The unconstrained model of GL led to an interpretation as a profit maximizing investment with the non-negativity of expenditures leading to a boundary (complementary slackness) condition for positive investment. Investment expenditures are increasing in Losses (GL, footnote 15 carries over). Although such a problem is developed in Farrow, he focused on a budget constrained problem in part because the initial organization of focus was Government. Such a constraint may apply in any organization with the un-constrained cost-minimization problem reproducing the unconstrained profit maximization function in these models.

Multiple sites and a budget constraint.

The budget constraint, if it is binding, effectively replaces the non-negativity constraint on expenditures in GL. In the unconstrained budget case (but with non-negativity of expenditure), the net benefit must simply exceed zero for the first expenditure. In the (binding) budget constrained case, the net benefits must exceed some positive value determined by the shadow price of the budget constraint. In that sense a binding budget constraint raises the quantitative threshold for investments to occur unless the budget is so large that the problem effectively becomes unconstrained, the GL model.

The concern for multiple sites, or in GL context—multiple information sets, leads to a requirement that the net value of investments (expected damages avoided plus investment cost) are to be equated across all the sites. In application, GL appear to follow such an approach for multiple information sets [Gordon and Leob, 2015].

3. Extensions of the GL model

Farrow investigated numerous extensions of the model which are not considered in GL. These extensions are motivated by issues such as:

1) Damage reduction: Investments may reduce not only the probability of an attack but the severity of the attack. When the defensive types of expenditures are identified separately, they should be invested in until the incremental return per dollar invested in the same across the categories. Recent cybersecurity concerns echo this sentiment with extending the focus beyond protecting access, to finding ways to limit damage internally but quickly finding a breach and also by external actions such as providing information to those who may have suffered damage. (see Farrow, models 1, 2)

2) Attacker diversion: Investments in defense by one organization may divert attacker effort to another site. If larger firms are better protected than smaller firms, whether in the defense industry or elsewhere; the probability of attack may increase at other sites that previously (see Farrow, model 1B)

3) Continuous asymmetric focus (advanced persistent attack?): Limitations on ability to defend a site or consequences of an attack may lead to optimal inequality of defense across sites and information sets as security may not be reducible to the level desired in the absence of such persistent attacks. (see Farrow, model 1A)

4) Shared filtering or defenses: to the extent that some defensive activities reduce damages at other sites through positive external public goods, the benefit of the investment includes the sum of the benefits. This may occur if government or the private sector provides central hacker detection at service providers which protects a number of sites and for which returns to scale may exist through minimal marginal cost for protecting more users. The public good technology may replace a private good technology where only one provider is protected and which may exhibit diminishing returns to scale. (Note that the potential for external public bads of a security breach are included in the base model such that external effects are part of the decision calculus). The social optimum is as presented in Farrow (see model 3) although game theoretic models for firm decision-making with free-riding or probability of a breach elsewhere in the network typically lead an individual firm to underinvest in shared information or individual defenses (Kunreuther and Heal, 2003; Gordon, Leob and Lucyshyn, 2003).

5) All hazards and/or false negative and false positive outcomes: As with most applications of uncertainty, there is often a chance of incorrectly applying a

defensive method; or of failing to apply a defensive method. For instance, as cybersecurity defense may move from “signature” identification of problematic sites or malware toward “behavioral” identification based on the actions of software; it may be that the probably of a “false positive” increases. A false positive in this case is when an actions is indeed appropriate, but the behaviorally based software identifies the action as inappropriate. There is a cost associated with such false positives, even if the focus is often on the costs of the false negatives--those actions that are indeed “bad” but are not identified as being bad. The investment model extends naturally to include the additional probabilities and costs associated with this broader set of actions. A similar multi-error concern may occur if a cyber-expenditure not only reduces the hazard of malware causing a cyber-physical system breach but also increases the probability of a false (positive) change to the cyber-physical system. (see Farrow, model 4)

These extensions of the cybersecurity model and their optimum conditions are summarized below. The reader is referred to Farrow (2007) for the formal model statement and further discussion.

Issue/model	**Summary of Socially Optimum Condition**
Allocating a Total Defensive Expenditure among Multiple Independent Sites	Equate the marginal social costs avoided (MSCA) across sites; possibly no protection at some sites
Advanced Persistent Threat: Technological or Behavioral Constraint on Probability or Cost Reduction	Technological or behavioral constraints can result in an optimal inequality among sites even where investment occurs
Allocation of Expenditures Across Damage and Probability-Reducing Activities	Equate the marginal social cost avoided of each type of expenditure where expenditures are positive (some may be below threshold)
Public Goods, Border control and	Invest until the sum of their marginal damage

Issue/model	Summary of Socially Optimum Condition
Positive Interdependencies	cost avoided equals the individual site MSCA
Site Interdependencies Due to Displacing the Probability of Attack.	Determine the net MSCA, net of probability increasing effects at other sites; sites of attack may be spread but social costs reduced
Multiple Sources of Probability (All Hazards) and Cost, as with false positives and false negatives from behavioral controls.	The form of the allocation decision is the same (e.g. equate MSCA), but all costs and probabilities should be taken into account

4. *Maximum cybersecurity investments*

Expected value optimization models generally bound investments to lie between zero and the expected value of losses based on the boundary conditions of the problem. An important part of GL was the additional step of considering specific functional forms for the security breach function, S. Their rather dramatic conclusion was that for the forms investigated, the optimal security investments would not exceed about 37 percent (1/*e*) of the initial expected loss. This conclusion has been the subject of additional research by Willemson (2006), Hausken (2006) and Baryshnikov (2012) indicating that while a significant class of functional forms (log convex with independent security effects of additional investments) follows the 1/*e* limit; other forms, such as linear, extend the full range of possible investment up to the value of the expected loss.

Recent work on log convexity and log concavity by Bagnoli and Bergstrom (2005) on probability functions inform the sensitivity of log convexity to functional form. Both v and S were identified as probability functions. Starting with S, the security breach probability; it represents the remaining probability of a breach (or break) after an investment z. As such, one interpretation is that it represents a reliability function (or exceedance function), the complement of a cumulative distribution function (Bagnoli and

Bergstrom, 2005). The initial probability of a breach, v, would then be the initial reliability which can range between zero and one and the difference between v and S is the increase in reliability or conversely the decrease in the probability of a breach. Bagnoli and Bergstrom (2005) investigate the log-concavity and convexity of numerous reliability functions (after investigating the implications of various transformations[4]). They catalog a large number of standardly used pdf, cdf and reliability functions. Log-concavity appears to dominate but some distributions are log-convex in their density functions and log-concave in other functions. The conclusion drawn by this author is that the functional form of the security breach function is an empirical question with numerous candidate forms implying the range of loss, in the expected value model, is from zero to the expected loss. Given the equivalent structure, quantification of any expenditure bounds in the Farrow model are expected to follow those of the GL model.

5. Relation between private and socially optimal investment

Optimal security expenditures increase as expected damages increase. (GL, 2002). Because GLLZ include external damages in their calculation of expected damages, they find that the socially optimal expenditure on cybersecurity increases as external damages increase. While GLLZ's model (p. 29) accounts for the external cost of a data breach, they demonstrate that a socially optimal expenditure on cybersecurity remains lower than the firm's expected private costs so long as the external costs remain below 180% of the private costs.[5] GLLZ call an external cost of more than 180% of the private costs

[4] Bagnoli and Bergstrom (2005, p. 452) show that log-convexity and log-concavity are preserved by negative transforms and truncation.

[5] For purposes of this section, the "private cost" of a data breach is the cost incurred by the directly attacked firm, and the "external cost" of a data breach is the cost incurred by the rest of society. This conception of "external cost" is narrower than the definition employed by economist Mark Cohen, who defines the "external cost" of a crime as "a cost imposed by one person onto another, where the latter person does not voluntarily accept the negative consequence." Cohen, Mark A. "Measuring the Costs and Benefits of Crime and Justice." Chapter in Volume 4 (pp. 263-316): Measurement and Analysis of Crime

"extremely large," (p.28), suggesting that a firm is ordinarily acting in society's best interest when it spends less on cybersecurity than its expected private cost. Campbell et al. (2003) and others have investigated the impact of cyber breaches on the stock market value of firms with varying results as to their significance and the effect if any is likely to be evolving over time. A micro-oriented legal approach is taken here to investigate the extent to which a firm suffering an attack is liable for damages to individuals or other firms. Although the initial firm is itself a victim, the incentives the firm faces to undertake defensive actions reflects in part the losses it may incur through legal action (or through insurance payments reflecting the potential for legal action among other factors).

A close look at several data breaches shows that an external cost of more than 180% of the private cost is not unusually large in some types of attacks[6]. Firms face the risk of many different sorts of data breaches. Some data breaches target consumers' financial information, others target a firm's proprietary information, while other breaches target critical infrastructure. Each type of breach presents its own potential for both private and external costs. Yet in many of the representative cases discussed below, external costs exceed 180% of private costs, and the legal system is unable to assign legal liability to data-storing firms in a manner that effectively internalizes the external costs of a data breach.

5.1 Personal identity theft

and Justice," Criminal Justice 2000. National Institute of Justice, July 2000, NCJ 182411; available at <http://www.ncjrs.org/criminal_justice2000/vol_4/04f.pdf>. What we're calling the "private cost" of a data breach would fall under Cohen's definition of an "external cost." For the purpose of this section, however, the "external cost" excludes the cost to the directly-attacked firm.

[6] Several cases were selected as being large and identifiable, others as setting important legal precedents, and others chosen at random from the data breach data base at the Privacy Rights Clearinghouse (www.privacyrights.org/)

Personal identity theft is one of the more visible breaches of cybersecurity when attackers may gain personal information from human resource departments, stored billing information or other databases. In PI cases, the perpetrator of a data breach generally seeks information about third-parties (e.g. customers) not information about the company whose servers are breached. Not surprisingly, these data breaches have the potential to create high levels of external costs.

When there are external costs of a personal information data breach, third-parties who lose money (including customers, their financial institutions, and their credit card companies) often sue the directly attacked firm. These plaintiffs have enjoyed varying degrees of success in the legal system. In some cases, plaintiffs have recovered some of their costs from legal settlements. When a plaintiff receives a settlement from the directly attacked firm, the external cost from the data breach is reduced and the private cost as defined by GL increases. In other cases, plaintiffs were not able to recover because they were unable to show standing, were prevented from recovering due to the economic loss doctrine, or faced some other legal or practical barrier to recovery. In these cases, the plaintiffs' external costs remain external costs and are not expected to enter the private benefit-cost calculation. Even where a consumer or financial institution was able to recover something, the external costs from a data breach can exceed 180% of the directly attacked firm's private costs.

The Heartland Payment Systems data breach is an example of third parties recovering some of the external costs of a data breach, but still losing more than 180% of the private cost. Heartland Payment Systems is a major payment processing company. In 2007, malware was implanted in Heartland's servers leading to the theft of 130 million customers' credit card. Heartland faced five lawsuits from third parties: a class-action suit from consumers; suits from Visa, MasterCard, and American Express, each of which filed the suits together with affiliated card-issuing financial institutions; and a suit from financial institutions that sued independently from the credit card companies. Four of the suits settled.

In settling the suits, Heartland partially internalized the external costs that the breach had imposed on financial institutions and consumers. Heartland paid a total of $105 million to VISA, MasterCard, American Express, and financial institutions that issued credit and debit cards through these companies. This appears to be far less than the full damages suffered by the plaintiffs. As Graves, Acquisti, and Christin (2014) showed, between 60% and 90% of compromised credit cards are reissued by financial institutions after a data breach. The same study showed that when second-order costs are considered, an issuer does not save money by not reissuing the cards. Therefore, a data breach costs the card issuer roughly the same amount regardless of whether the issuer reissues the cards or not. The cost of reissuing credit cards ranges from $2.70 to $11 per card, while the cost of reissuing debit cards ranges from $2.99 to $12.75 per card.[7] The cost depends on the size of the financial institution; large banks can reissue for less than smaller community banks and credit unions. The theft of 130 million records from Heartland, meaning that the credit card companies and affiliated financial institutions lost between $350 million (assuming a loss of $2.70/card) and $1.7 billion (assuming $12.75). The $105 million in settlements that the financial institutions received in settlements was a fraction of their external costs. Assuming that the entire settlement went to the financial institutions (in reality, some of it was consumed in legal expenses) the financial institutions' uncompensated losses would be somewhere between $250 million and $1.6 billion.

Consumers also lost money in the Heartland Security breach, and filed a class-action suit against Heartland. Heartland settled the suit for $3 million. Under the terms of the

[7] Crosman (2014).

settlement, consumers were eligible for up to $175 in compensation if they could show, by a preponderance of the evidence, that they had lost time or money cancelling a credit card or as the result of unauthorized credit cards. Consumers were also eligible for up to $10,000 if they could show that they had suffered identity theft as a result of the breach. Only 11 consumers successfully qualified for relief, and they received a combined total of just $1,925. The rest of the $3 million settlement was spent on legal fees, administrative costs, and a *cy pres* payment to a non-profit focused on information security.[8] While the settlement increased Heartland's private costs by $3 million, it only compensated consumers' external costs by $1,925. It is inconceivable that $1,925 represents the socially optimal expenditure to avoid external costs for the 130 million consumers.

Heartland has estimated its private cost from the data breach as $140 million, a figure that includes the $108 million in settlements ($105 million to the financial institutions, and $3 million to the consumers), plus legal fees and other expenses. (Silver-Greenberg and Schwartz, 2012). It is challenging to put an exact figure on the external losses, but it is clear that they are greater than $250 million, which would be 180% of Heartland's private costs. As shown above, the external cost to financial institutions—after subtracting the $108 million in legal settlements—is likely between $250 million and $1.6 billion. (Without the $108 million in settlements, the financial institutions would have lost between $350 million and $1.7 billion.) Since the external cost to consumers is not negligible (though hard to value) the combined external cost to consumers and financial institutions is certainly greater than 180% of Heartland's private costs.

While external costs in the Heartland data breach substantially exceeded private costs, the discrepancy would have been even greater had Heartland's legal settlements not

[8] *In Re: Heartland Payment Systems, Inc. Customer Data Security Breach Litigation*, 851 F. Supp. 2d 1040 (S.D. Tex. 2012).

decreased the external costs and increased the private costs. In other major data breaches, courts have thrown out lawsuits against directly attacked firms. In these data breaches, external costs as a percentage of private costs can be even greater than in Heartland.

Third parties who incur costs as the result of a data breach have struggled to show standing (a sufficient legal basis on which to file suit) since the Supreme Court's 2013 decision in *Clapper v. Amnesty International*. (133 S. Ct. 1138 (2013)). A requirement for standing in a federal court is that a plaintiff must have suffered an "injury in fact," which is "concrete and particularized" instead of "merely speculative." (*Lujan v. Defenders of Wildlife*, 504 U.S. 555, 560-61 (1992) (citations omitted)). In *Clapper*, the Court held that the "injury in fact" must be something that either already happened or is "certainly impending," not something that might happen in the future. (133 S. Ct. at 1151.) Furthermore, a plaintiff cannot "manufacture" standing by "incurr[ing] certain costs as a reasonable reaction to a risk of [future] harm." (*Id.*) In a data breach case, a consumer will often wish to sue before becoming the victim of identity theft. Similarly, a financial institution may wish to reissue credit cards before the credit or debit cards are misused. Under the rule announced in *Clapper*, the legal system may be unable to reimburse third-party data breach victims for these costs likely increasing the external component of losses.

The Zappos data breach, which occurred in 2012, is an example of a data breach in which the legal system has done little to turn external costs into private costs. The perpetrators of the breach stole 24 million customers' names, email addresses, billing and shipping addresses, phone numbers, the last four digits of payment card numbers, and encrypted passwords. No credit card numbers were compromised, so financial institutions did not sue. A class-action suit from Zappos consumers alleged that they faced an increased risk of identity theft. In addition to facing an increased risk of identity theft, the consumers asked to be compensated for credit monitoring purchases that they had made. The United States District Court for the District of Nevada cited *Clapper* in dismissing the suit, holding that the consumers did not suffer the sort of injury that could confer standing in federal court. Judge Robert C. Jones wrote that his court "realizes that [dismissing the suit] is a frustrating result where Plaintiffs' fears of identity theft and fraud are rational, and it recognizes that purchasing monitoring services is a responsible response to a data breach. Nevertheless, costs incurred to prevent future harm is not enough to confer standing, even when such efforts are sensible." *In re Zappos.com, Inc., Customer Data Sec. Breach Litig.*, No. 3:12-cv-00325-RCJ-VPC, (D. Nev.) 2015. While the consumer litigation was dismissed,[9] Zappos settled a separate lawsuit with nine state attorneys general for $106,000. But aside from this payment, none of the data breach's external costs were converted to private costs.

Third parties who suffer external costs in personal information data breaches can also be prevented from recovering by the economic loss rule. In some--but not all--states, the economic loss rule prevents plaintiffs from recovering for negligence when they have only lost money as a result of the alleged negligence. In these states, financial disputes can only be adjudicated through contract law. If a plaintiff cannot show that the defendant breached a contract, then she cannot recover for negligence. The economic loss rule prevented recovery for employees of the University of Pittsburgh Medical Center. These employees were unable to sue their employer for negligence when the hospital lost their personal information when a court found that the employees' only losses had been financial. (Willett, 2015)

5.2 Intellectual Property Theft

[9] The case was dismissed without prejudice, allowing the consumers to submit the suit again if they could show an injury in fact.

In cases of intellectual property theft, external losses but not necessarily "private" losses are likely to be low in the short-term, well below the 180% level that GL call "extremely large." Some intellectual property data breaches are committed by insiders: employees--often departing ones--who steal large numbers of records from their employer. These insider breaches of intellectual property are likely to have considerable private costs, although companies that are the victims of these incidents have been fairly successful at using the legal system to minimize those losses. In contrast, American companies impacted by foreign intellectual property attacks--like the Chinese military's Unit 61398--have generally been unsuccessful at using the legal system to ameliorate their private costs. Yet in either case, there is little external cost. For example, when a departing bank employee steals a list of customers who have applied for loans, the customer loses little. In fact, the customer may actually benefit from the data breach. If, for example, the departing employee uses the stolen information to offer the customer a lower interest rate, the external cost of the data breach could actually be negative. In the long run, however, this sort of intellectual property theft could create some external cost by increasing the cost of information security and decreasing the overall competitiveness of American firms, an external cost to the country.

The 2011 litigation over an intellectual property data breach at Huntington National Bank illustrates the fact that these sort of data breaches tend to have low (or no) external costs, and private costs that are generally recoverable through litigation. (*The Huntington National Bank v. Kokoska et al,* Docket No. 1:11-cv-00063 (N.D. W. Va. Apr 25, 2011)). The Huntington breach occurred when employees of Huntington National Bank—loan officers and their administrative assistants—left Huntington for MVB Bank, a rival financial institution which (like Huntington) issued mortgages. Huntington accused the employees of downloading records from 200 loan applications before leaving Huntington for MVB. Huntington alleged that as a result of the ex-employees' conduct, they suffered damage to their reputation and goodwill in the marketplace, lost customers, and lost rights of exclusive possession in their intellectual property. The United States District Court for the Northern District of West Virginia issued a temporary injunction preventing the defendants from using any proprietary information obtained from Huntington National Bank for any business purpose. After the court issued its injunction, the two sides settled. The settlement obligated the ex-employees to refrain from using the proprietary information, to return the documents from Huntington National Bank, and not to contact Huntington customers with whom the ex-employees worked. As stated above, this sort of conduct creates little short-term potential for external costs, since the 200 customers probably would have benefited from MVB offering a more attractive mortgage than the one offered to them by Huntington. This sort of data breach also offers little potential for private costs, since Huntington obtained a consent decree which prevented MVB from accessing any records the employees might have taken from Huntington, or even contacting any customers the employees worked with at Huntington.

It is harder for American companies that suffer intellectual property data breaches from non-U.S. perpetrators to use the legal system to recover their private costs. Westinghouse, an American nuclear power company, suffered several data breaches when hackers from Unit 61398 of China's People's Liberation Army (PLA) broke into the company's network and stole information about the company's strategy for negotiating with a Chinese counterpart. (Schmidt and Sanger, 2014). As with the data breach at Huntington National Bank, this theft created almost no short-term external costs. In fact, third-parties might have even benefitted if Westinghouse's Chinese competitors were able to use the stolen information to generate power more inexpensively. The breach likely did create private costs in GL terms, especially if Westinghouse ended up at a disadvantage in its negotiations with the Chinese firm, or if the company lost market-share to a Chinese competitor with lower costs. A grand jury at the United States District Court for the Western District of Pennsylvania indicted five officers of Unit 61398 for their role in data breaches against Westinghouse and other American firms. This indictment is unlikely to reduce Westinghouse's private costs, since China will almost certainly not extradite the indicted officers. In fact, Westinghouse's private costs could actually increase as a result of its stepping forward and identifying itself as a victim of Unit 61398, since China could retaliate against the company's Chinese interests to punish the company for accusing the PLA unit. (Schmidt and Sanger, 2014).

Speculation exists that longer term, macroeconomic external effects may occur with large scale intellectual property thefts (Andrijcic and Horowitz, 2006). A loss of macro-level comparative advantage, for instance in technologically advanced areas, for the United States could affect the welfare of the US labor force in the absence of the perfectly competitive labor and international markets as well as affecting US based shareholders.

5.3 Critical Infrastructure Cyberattacks

A cyberattack that targeted critical infrastructure could create external costs that greatly exceed private costs. The Critical Infrastructures Protection Act of 2001, a federal law that is part of the USA PATRIOT Act, defines critical infrastructure as "systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters."[10] The Obama Administration has designated 16 sectors of the economy as critical infrastructure.[11, 12]

Like data breaches that target personal information or intellectual property, cyberattacks on critical infrastructure involve a perpetrator exercising control over a computer network to which they do not have authorized access or authority. There are, however, important differences between cyberattacks on CI and data breaches that target PI or IP. Perpetrators of PI and IP data breaches primarily attempt to steal information for economic benefit. Perpetrators of CI cyberattacks, on the other hand, are typically motivated by strategic or political goals.

The United States has never experienced a successful, large-scale cyberattack on critical infrastructure, but other countries have. The Russian government is widely believed[13] to have conducted cyberattacks on critical infrastructure in both Estonia and Georgia. Estonia experienced cyberattacks in 2007 as the Baltic nation was engaged in a dispute with Russia about a monument to Soviet war dead. Many Estonians viewed the monument as a reminder of the brutal Soviet occupation of their country, while Russians viewed the memorial as expressing respect to Soviet soldiers who fell in battle during World War II. (Tanner, 2007). On April 26, 2007--one day before the Estonian

[10] This can be found at 42 U.S.C. § 5195c(e).

[11] Presidential Policy Directive 21, available at https://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil

[12] The 16 sectors are: chemical; commercial facilities; communications; critical manufacturing, dams, defense industrial base, emergency services, energy, financial services, food and agriculture, government facilities, healthcare and public health, information technology, nuclear reactors, materials, and waste; transportation systems; and water and wastewater systems

[13] Russia has never claimed responsibility for the attacks on Georgia, although (as explained below) the cyberattacks appeared to be coordinated with the movements of conventional Russian military forces. A Russian official did, however, eventually confirm the participation of a Russian government-sponsored youth group in the denial-of-service attacks on Estonian websites. Miller, C. Russia Confirms Involvement with Estonia DDOS Attacks., *SC Magazine*, March 12, 2009. Available at: <http://www.scmagazine.com/russia-confirms-involvement-with-estonia-ddos-attacks/article/128737/>

government was to move the statue from the center of Tallin to a military cemetery--many Estonian websites experienced denial-of-service attacks. The websites of Estonian banks, government agencies, and media outlets were overwhelmed with traffic, causing them to crash. On April 27, the statue was moved as planned, but the attacks continued, reaching their peak on May 9. Desperate to stop the attacks, the Estonian government took the comprehensive action of blocking all internet traffic from outside the country. By May 19, 2007, the attacks ended.

Georgia experienced similar denial-of-service attacks during its brief war with Russia in the summer of 2008. Hackers targeted the websites of Georgia's national and local governments, as well as news websites. (Hollis, 2011). The National Bank of Georgia's site was also briefly targeted. (Markoff, 2008). The cyberattacks appeared to be coordinated with Russia's military operations: the attacks escalated as Russian troops crossed the Georgian border, and targeted websites of specific regions of Georgia as Russia bombed those areas. The cyberattacks were also coordinated with Russia's strategic objections in the war. Just as Russia damaged areas around Georgia's strategic Baku-Ceyhan oil pipeline but did not actually destroy the pipeline, the Russian hackers targeted government and media websites, but declined to target Georgia's electric grid or attack the National Bank's site in a more sustained manner. The impression is that Russia could have done worse damage later if it wanted.

Both of these cyberattacks demonstrated that a cyberattack on critical infrastructure can produce external costs that are significantly higher than the private costs. A leaked American diplomatic cable estimated that Hansabank, the bank whose website was most disabled by the cyberattack, spent €10 million (roughly $14 million, in 2007 dollars) as a result of the attack. (Keizer, 2010). €10 million is a suspiciously round number, and reflects the challenge of reaching an accurate estimate. The *New York Times* reported at the end of May 2007 (after the last wave of cyberattacks) that Hansabank had spent $1 million so far fending off the so-called “bots” that were overwhelming its servers. (Landler and Markoff, 2007). Yet both these estimates only reflect what Hansabank spent fending off the attackers. The bank also lost significant business during the three weeks when its website was either nonoperational or closed off to traffic from outside of Estonia.

Assessing the external cost of the cyberattacks on Estonian banks and other websites is more difficult. Estonia is one of the most technologically sophisticated countries in the world, and by 2007, over 96% of Estonian banking transactions took place online. (Richards, 2009). When Estonian bank websites were either non-operational or closed off to foreign traffic, the Estonian economy was significantly constrained. Yet the damage to the Estonian economy was not catastrophic, and not as bad as it could have been had the cyberattack been marginally more aggressive. The same leaked American diplomatic cable noted that the damage to the Estonian economy could have been significantly worse had the “second wave” of attacks targeted the “poorly-defended” websites of the logistics firms that transport food and gasoline around the country. (Hõbemägi, 2010). One indication of the cyberattacks’ high external cost is the lengths to which the Estonian government has gone to prevent another wave of cyberattacks. Since 2007, Estonia has become a world leader in defending against cyberwarfare, and has played a crucial role in NATO and EU efforts to respond to cyberattacks. In 2008, largely in response to the attacks on Estonia, NATO created a Brussels-based Cyber Defense Management Authority and an Estonia-based Cooperative Cyber Defense Centre of Excellence. (Herzog, 2011). In 2009 the Estonian government created a Cyber Security Council, which is responsible for coordinating the government’s efforts to protect Estonia from online threats to its national and economic security. In 2014, the

Council proposed a four-year strategy for countering cyberattacks that will cost almost €16 million.[14]

Because the cyberattacks against Georgia occurred in the context of a conventional war and were coordinated with war objectives, it is hard to isolate the private and external costs specific to the cyberattack. David Hollis explains how the cyberattacks were integrated into Russia's war strategy:

> Russian-oriented hackers/militia took out news and local government web sites specifically in the areas that the Russian military intended to attack in the ground and air domains. The Federal and local Georgian governments, military, and local news agencies were unable to communicate with Georgian citizens that were directly affected by the fighting. …[The cyberattacks] created panic and confusion in the local populace, further hindering Georgian military response. (Hollis, 2011).

The cyberattacks didn't just cost Georgia money; they weakened Georgia's ability to fight. The external cost was military, not just financial.

Russia's goal in launching the cyberattacks against Georgia and Estonia was presumably strategic, not economic. Russia sought to assert its power throughout the former Soviet Union. To accomplish that goal, Russia did not need to cause massive damage to the economy of Georgia or Estonia; it just needed to show what it was capable of doing. Simply by demonstrating the power to disable crucial elements of its rivals' economies, Russia increased its ability to deter its neighbors from defying it. Today, it is likely that when countries in the former Soviet Union (including Estonia and Georgia) consider taking some action that would displease the Kremlin, they consider the possibility that a

[14] An English-language copy of the report is available at <https://www.mkm.ee/sites/default/files/cyber_security_strategy_2014-2017_public_version.pdf>

cyberattack could paralyze their economy. Some countries may minimize their risk of a Russian cyberattack by making political concessions, or refraining from adopting policies likely to anger Russia. This is an external cost as well.

If the United States experienced a cyberattack on critical infrastructure, the external losses could greatly exceed the private costs. The costs would, of course, depend on what sort of critical infrastructure was attacked. One threat that looms particularly large is the threat of a cyberattack on the electric grid. In 2014, NSA Director Admiral Mike Rogers warned Congress that China and at least one other country had the potential to disrupt large sections of the American power grid. Rogers warned that other countries were already conducting “reconnaissance” to figure out how American electric networks operated, and how they could be shut down online. (Crawford, 2014). According to a 2015 report from the University of Cambridge Centre for Risk Studies and the Lloyd's of London insurance market, a cyberattack that targeted electricity supply could create huge costs to the American economy. The report models several different attacks. In each one, the external cost of a cyberattack to the economy as a whole greatly exceeds the costs borne by individual electric companies. For example, in a scenario where the U.S. economy loses $1.024 trillion as the result of a catastrophic interruption in the power supply, power companies would only lose $4.21 billion in lost revenue. In a scenario where the economy loses $243 billion, the power companies would only lose $1.15 billion. (Lloyd’s, 2015). Because a catastrophic power outage would be extremely disruptive to the economy as a whole, external costs could exceed private costs by two orders of magnitude.

The American legal system may be able to transfer some of the external cost of a critical infrastructure data breach onto the company that suffered the attack. It is difficult, but sometimes possible, to successfully sue a power company for failing to prevent a power outage. In nearly every state of the country, power companies are allowed to limit their liability for an electric outage to cases of "gross negligence," as opposed to the sort of ordinary negligence for which most companies are held liable.[15] (Liptak, 2003). In New York, for example, a public utility can limit its liability to "gross negligence" or "willful misconduct," and exempt itself from being sued for ordinary negligence. (*Food Pageant, Inc. v. Consol. Edison Co.*, 429 N.E.2d 738, 740 (1981)). In Maryland, a public utility is allowed to limit its liability to "willful neglect" or "willful default," which is a similar standard. (*Singer Co., Link Simulation Sys. Div. v. Baltimore Gas & Elec. Co.*, 558 A.2d 419, 428 (1989).)

Customers who experience losses as the result of a power outage sometimes are able to show "gross negligence" and recover from electric companies. After the 1977 blackout in New York City, for example, a grocery store successfully recovered from the Con Edison power company when it showed that Con Edison failed to take a number of safety precautions that could have prevented or limited the damage caused by the blackout. *Food Pageant, Inc. v. Consol. Edison Co.*, 429 N.E.2d 738 (1981). Similarly, after Hurricane Sandy, Con Ed reached a settlement with homeowners who lost power in the storm. The power company paid $17 million, and agreed to cancel a $40 million rate increase. (Frankel, 2012). It is possible that a third-party victim of a cyberattack on an electric company could sue the electric company, effectively turning the third-party's external costs into the company's private costs. The electricity sector is subject to binding cybersecurity regulation. (Zhang, 2013). If an electric company was grossly negligent in failing to meet these cybersecurity standards, it could face liability which would convert its customers' external costs into company private costs.

[15] *See Garrison v. Pac. Nw. Bell*, 608 P.2d 1206, 1211 (1980) ("Courts are virtually unanimous that provisions limiting a public utility's liability are valid so long as they do not purport to grant immunity or limit liability for gross negligence.")

Despite this thin reed of legal opportunity, third-parties who suffer losses in a critical infrastructure cyberattack are unlikely to recover much of their losses by filing a lawsuit for gross negligence. Occasional successes notwithstanding, it is generally very difficult to show that a public utility was grossly negligent. (Wei, Debaise, and Bray, 2003). Furthermore, in a truly catastrophic critical infrastructure cyberattack, an affected company may not have enough money to compensate all the third-party victims. Recognizing that the tort system was a poor means of protecting consumers from cyberattacks, Congress passed the Support Anti-Terrorism by Fostering Effective Technologies Act (SAFETY Act) of 2002. The SAFETY Act protects companies from liability in the event of a physical or cyberattack, on the condition that the companies employ technology that the Department of Homeland Security finds can be effective at preventing the attack from occurring. Technologies can be certified by the Department of Homeland Security if they comply with the best practices in counterterrorism, and if certified, they are protected from legal liability. (Venable, 2014). The SAFETY Act seems to assume that government and the legal system are better situated to prevent attacks in the first place than to adjudicate liability in the event of an attack.

6. Large Losses and Risk Aversion

GL explicitly focused on small losses and used a risk neutral model (GL, 2002, p. 440-441). Some, although not all, cybersecurity breaches could entail large losses such that private or public decision-makers may be risk averse. For instance, a loss of intellectual property that is the primary asset of a firm, or a breach that endangers a linked cyber-physical system such as the power grid, water, or some transportation modes. In such instances private sector decision-makers are routinely modeled as being risk averse instead of risk neutral. In general, risk aversion implies a willingness to invest to avoid a risky outcome (for instance, a security breach) that exceeds the expected loss in contrast with risk neutrality, as above, where the investment is limited by the expected value (Eeckhoudt, Gollier and Schlesinger, 2005). While this result also occurs in the cybersecurity literature, it is also demonstrated that the willingness to invest has an upper bound of conditional loss even with risk aversion (Huang, Hu and Bahara, 2008) while investments can be less than the expected loss depending on the nature of the asset being threatened, for instance, it if is "irreplaceable" (Cook and Graham, 1977).

Economists have considered whether national level, public decision-makers "should" be risk neutral or risk averse. Influential work by Arrow and Lind presented a model why such decision-makers "should" be risk neutral when, in fact, public decisions are often accepted where the costs of control appear larger than the expected value implying risk averse decision-making (Lucas, 2014). Stewart, Ellingwood and Mueller (2011) have contrasted expected value and risk averse decision-making for an illustrative security investment, and in Stewart and Mueller (2013) applied that analysis to the Transportation Security Agency using a discrete expected loss model equivalent to that of GL and Farrow. Their conceptual results (using a specific functional form for utility) are informative and reproduced in Figure 1. That figure is to be read such that some utility level equates a policy of doing nothing with a policy of the specified security investment. At that utility level, a risk neutral decision-maker would at most spend the expected loss as read from the straight line connecting utility and cost (the loss). All the lines to the right of the straight line represent varying levels of risk aversion and the horizontal "cost" difference between them represents the amount in excess of the expected loss which a risk averse decision-maker is willing to pay given the functional forms used. Such sums can be significant. In two examples discussed in the text (Stewart, Ellingwood, and Mueller, 2011, p. 381, 382), the optimal defensive investment exceed the expected loss by factors of 2.97 and 125 respectively noting that with small probabilities the expected loss can be a small fraction of the conditional loss.

Figure 1: Utility and Risk Neutral and Risk Averse Costs

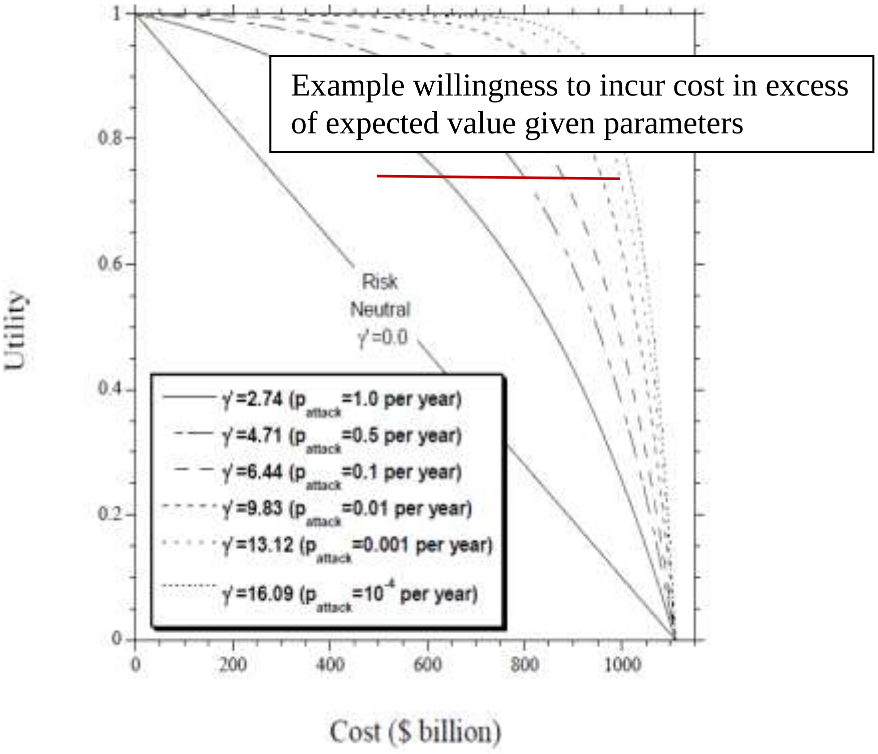

Example willingness to incur cost in excess of expected value given parameters

Source: Stewart, Ellingwood and Mueller, 2011

The specific point in the context of the GL two outcome model is that expenditures with risk aversion can significantly exceed expected losses in the discrete outcome model.

However, the result for continuous state models are somewhat more ambiguous. While the general point remains that a risk averse decision-maker will expend more than a risk neutral decision-maker, the difference need not be large. In a continuous state setting, damage functions and probability are both varying. If damages are rising faster than the probability is declining, it is possible that no mean value exists. In contrast, for a concave damage function investigated for flooding; the mean values using expected value, a risk averse utility function, and a weighting based on cumulative prospect theory were not substantially different (Farrow and Scott, 2012). In part, the large losses, where risk aversion can have a significant effect, were associated with sufficiently small probabilities that there was not a large distinction between the mean values using a risk neutral or a risk averse utility function. Like the finding for the security breach function, the importance of risk aversion is an empirical question involving the structural form of damages, probabilities and the utility function.

7. Conclusion

GL thoroughly developed an influential model of cyber-security investment with important implications on the size of cyber-security investment. As with the history of economic modeling, if one is lucky enough to get clear initial results; then such results often get extended and qualified. Such is the main theme of this note. The key issues and conclusions identified here are:

- Consideration of externalities is a primary concern for government policy. Although GLLZ imply that externalities are unlikely to be large; evidence from past data breaches and cyberattacks suggests that externalities may be large, and that the legal system often fails to significantly internalize external costs by allocating risks to the original target. Such allocation may or may not be viewed as fair. This conclusion is relatively clear for personal identity theft and infrastructure attacks while the external effect of intellectual property attacks is less well documented in legal findings. These findings can imply significantly larger cybersecurity expenditures than those based solely on internalized, private sector damages.

- While an *expected* value (risk neutral) decision-maker will not spend more than expected losses, the empirical functional form affects how much less, if any, than the expected value of damages is invested.
- When risk neutral investment considerations include: a) multiple datasets or sites, b) public defenses that protect many sites (as opposed to sharing of information), or c) an ability to invest in reducing damages, then the optimal (constrained) level of cyber investment changes and typically requires economic security based on the equivalency of marginal external social cost across conditions to the extent technologically possible.
- When additional modeling of the (dis) utility of uncertainty is added, whether in the form of irreversible losses or risk aversion; models commonly exist in which investments exceed expected losses.

The overall conclusion drawn here is that while there may be useful rules of thumb for decision-making, those behavioral rules may be incorrect depending on the empirical context of a specific problem including the behavior of attackers. To the extent that government policies and investments tend to focus on cyber issues with larger externalities, make use of public defenses, or be subject to advanced persistent threats, then the socially optimal investment decisions may be larger and differ significantly from those of the private sector.

References

Andrijcic, E. and B. Horowitz, 2006. A Macro-Economic Framework for Evaluation of Cyber Security Risks Related to Protection of Intellectual Property *Risk Analysis,* 26(4):907-923.

Bagnoli, M. and T. Bergstrom, 2005. Log-Concave Probability and Its Applications, *Economic Theory* 26(2): 445–469

Baryshnikov, Y., 2012. IT Security Investment and Gordon-Loeb's 1/e Rule, *Proceedings of the 11th Workshop on the Economics of Information Security (WEIS).I*

Campbell, K., Gordon, L. A., Loeb, M. P., & Zhou, L., 2003. The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security*, *11*(3), 431-448.

Cook P., and Graham, D., 1977. The Demand for Insurance and Protection: A Case of Irreplaceable Commodities. *Quarterly Journal of Economics*, 92:143-156.

Crawford, J., 2014. The U.S. government thinks China could take down the power grid, *CNN.com*, November 21, 2014. Available at <http://www.cnn.com/2014/11/20/politics/nsa-china-power-grid/>

Crosman, P., 2014. How Much Do Data Breaches Cost? Two Studies Attempt a Tally, *American Banker*, availiable at <http://www.americanbanker.com/issues/179_176/how-much-do-data-breaches-cost-two-studies-attempt-a-tally-1069893-1.html>

Eeckhoudt, L., C. Gollier, and H. Schlesinger, 2005. *Economic and Financial Decisions Under Risk*, Princeton University Press, Princeton, NJ.

Farrow, S., 2007. The Economics of Homeland Security Expenditures: Foundational Expected Cost-Effectiveness Approaches, *Contemporary Economic Policy*, 25(1):14-26.

Frankel, A., 2012. Can customers sue power companies for outages? Yes, but it's hard to win, *Reuters.com*, November 9, 2012. Available at <http://blogs.reuters.com/alison-frankel/2012/11/09/can-customers-sue-power-companies-for-outages-yes-but-its-hard-to-win/>

Gordon, L. and M. Loeb, 2002. The Economics of Information Security Investment, *ACM Transactions on Information and System Security*, 5(4):438-457.

Gordon, L. and M. Loeb, 2011. You May Be Fighting the Wrong Security Battles. *Wall Street Journal*, September 26.

Gordon, L., M. Loeb, and W. Lucyshyn, W., 2003. Sharing information on computer systems security: An economic analysis, *Journal of Accounting and Public Policy*, *22*(6), 461-485.

Gordon, L., M. Loeb, W. Lucyshyn and L. Zhou, 2015. Externalities and the Magnitude of Cyber Security Underinvestment by Private Sector Firms: A Modification of the Gordon-Loeb Model, *Journal of Information Security*, 6:24-30.

Graves, J., Acquiti, A., and Christin N., 2014. Should Payment Card Issuers Reissue Cards in Response to a Data Breach?, *WEIS: Workshop on the Economics of Information Security*, available at <http://www.econinfosec.org/archive/weis2014/papers/GravesAcquistiChristin-WEIS2014.pdf>

Hausken, K. 2006. Returns to Information Security Investment: The Effect of Alternative Information Security Breach Functions on Optimal Investment and Sensiviity to Vulnerability, *Information Systems Frontiers*, 8(5):338-349.

Herzog, S. 2011. Revisiting the Estonian Cyber Attacks: Digital Threats and Multinational Responses, *Journal of Strategic Security*, 4, no. 2 (2011): 49-60. Available at <http://scholarcommons.usf.edu/cgi/viewcontent.cgi?article=1105&context=jss>

Hõbemägi, T., 2010. Price of cyberattacks to Hansabank: 10 million euros, *Baltic Business News*, August 12, 2010. Available at <http://balticbusinessnews.com/article/2010/12/08/Price-of-cyberattacks-to-Hansabank-10-million-euros>

Hollis, D., 2011., Cyberware Case Study: Georgia 2008, *Small Wars Journal*, January 6, 2011. Available at <http://smallwarsjournal.com/blog/journal/docs-temp/639-hollis.pdf>

Huang, C.D., Q. Hu and R. S. Behara, 2008. An economic analysis of the optimal information security investment in the case of a risk-averse firm, *International Journal of Production Economics*, 114:793– 804.

Keizer, G., 2010. Estonia blamed Russia for backing 2007 cyberattacks, says leaked cable, *Computer World*, December 9, 2010. Available at <http://www.computerworld.com/article/2511704/vertical-it/estonia-blamed-russia-for-backing-2007-cyberattacks--says-leaked-cable.html>

Kunreuther, H., and Heal, G., 2003. Interdependent Security, *Journal of Risk and Uncertainty,* 26 (2-3):231-49.

Landler, M., and Markoff, J., 2007. Digital Fears Emerge After Data Siege in Estonia, *The New York Times*, May 29, 2007. Available at <http://www.nytimes.com/2007/05/29/technology/29estonia.html?pagewanted=all>

Liptak, A., 2003. THE BLACKOUT OF 2003: LAWSUITS; Plaintiffs to Face Hurdles Proving Liability, *The New York Times*, August 15, 2003. Available at <http://www.nytimes.com/2003/08/15/us/the-blackout-of-2003-lawsuits-plaintiffs-to-face-hurdles-proving-liability.html>

Lloyd’s of London, 2015. Business Blackout: The insurance implications of a cyber attack on the US power grid, Lloyd’s Emerging Risk Report - 2015. Available at <https://www.lloyds.com/~/media/files/news%20and%20insight/risk%20insight/2015/business%20blackout/business%20blackout20150708.pdf>

Lucas, D., 2014. Rebutting Arrow and Lind: why governments should use market rates for discounting, *Journal of Natural Resources Policy Research*, 6(1):85-91.

Markoff, J., 2008. Before the Gunfire, Cyberattacks, *The New York Times*, August 13, 2008. Available at <http://www.nytimes.com/2008/08/13/technology/13cyber.html?_r=0>

Richards, J., 2009. Denial-of-Service: The Estonian Cyberwar and Its Implications for U.S. National Security, *International Affairs Review*, Vol. 18, No. 2. (2009). Available at <http://www.iar-gwu.org/node/65>

Schmidt, M., and Sanger, D., 2014. 5 in China Army Face U.S. Charges of Cyberattacks. *The New York Times*, May 19, 2014. Available at <http://www.nytimes.com/2014/05/20/us/us-to-charge-chinese-workers-with-cyberspying.html>

Silver-Greenberg, J., and Schwartz, N. MasterCard and Visa Investigate Data Breach. *The New York Times*, March 31, 2012. Available at <http://www.nytimes.com/2012/03/31/business/mastercard-and-visa-look-into-possible-attack.html?_r=0>

Stewart, M, B. Ellingwood and J. Mueller, 2011. Homeland Security: A Case Study in Risk Aversion for Public Decision Making, *Int. J. Risk Assessment and Management,* 15(5/6): 367-386.

Stewart, M. and J. Mueller, 2013. Aviation Security, Risk Assessment, and Risk Aversion for Public Decisionmaking, *Journal of Policy Analysis and Management*, 32(3): 615–633.

Tanner, J., 2007. Estonia Moves Soviet Statue to Cemetery, *The Associated Press*, April 30, 2007.
Available at <http://www.washingtonpost.com/wp-dyn/content/article/2007/04/30/AR2007043000478.html>

Venable, LLP., 2014. The SAFETY Act: Providing Critical Liability Protections for Cyber and Physical Security Efforts, Available at <https://www.venable.com/files/Publication/6c0b031e-c2c5-4029-9ac7-13cb1d8c0d07/Presentation/PublicationAttachment/e81d24a3-fc57-4ece-8e1f-179418baf994/The_SAFETY_Act_Providing_Critical_Liability_Protections_for_Cyber_and_Physical_Securi.pdf>

Wei, L., Debaise, C., and Bray, C., 2003. Blackout Exposes Power Companies To Potential Litigation, *Dow Jones Newswires New York*, August 18, 2003. Available at <http://www.oandb.com/blackoutexposes.html>

Willett, B., 2015. Employees Can’t Sue Hospital for Negligence, Breach of Contract, After Personal Data Breach, *Reed Smith Technology Law Dispatch*, June 12, 2015. Available at <http://www.technologylawdispatch.com/2015/06/privacy-data-protection/employees-cant-sue-hospital-for-negligence-breach-of-contract-after-personal-data-breach/>

Williamson, J., 2006. On the Gordon and Leob Model for Information Security Investment, *Proceedings of the 5th Workshop on the Economics of Information Security (WEIS).*

Zhang, Z., 2013. *Cybersecurity Policy for the Electricity Sector: The First Step to Protecting Our Critical Infrastructure from Cyber Threats,* Boston University Journal of Science and Technology Law, Vol. 19 pp. 319-366.