Data Breaches and Identity Theft: When is Mandatory Disclosure Optimal?

Sasha Romanosky,[*] Richard Sharp,[**] Alessandro Acquisti [*]
[*] Heinz College of Information Systems and Public Policy, Carnegie Mellon University
[**] Department of Mathematical Sciences, Carnegie Mellon University

DRAFT

Abstract

In order to reduce identity theft and consumer loss caused by data breaches, many U.S. states have enacted laws requiring firms to notify individuals when their personal information has been stolen or lost. The effect of these disclosure laws has yet to be rigorously tested, and some claim that they only serve to burden firms and consumers with unnecessary costs. Leveraging the economic analysis of accident law, we examine whether mandatory disclosure policies can ever reduce overall social costs by inducing firms and consumers to take optimal care. Using both analytical and numerical modeling, we show that even though firm costs will be higher under disclosure regimes, firms can be induced to increase their investment in care, which may lower social costs. Moreover, disclosure can induce consumers to increase their level of care, thus lowering their total costs. Finally, we find that the change in social costs are typically increasing in disclosure 'tax' (costs imposed on firms due to disclosure laws) and decreasing in consumer redress (compensation paid to consumers by firms). However, when firms compensate consumers for only a small amount of loss, some disclosure tax may be necessary to optimally reduce social costs.

Keywords

Information disclosure, economic analysis of tort law, data breach, security breach notification, identity theft, analytical modeling

Acknowledgements

We would like to thank CyLab at Carnegie Mellon University for their generous support under grant DAAD19-02-1-0389 from the Army Research Office.

1. Introduction

This paper is concerned with data breaches and resulting consumer privacy harms, such as identity theft. Data breaches occur when personal consumer information is lost or stolen, and can result in the loss of hundreds or millions of records (e.g., local schools or small retail stores; TJX or Heartland). They can occur from the improper disposal of documents containing personal information, from the loss of a laptop or thumb-drive, or when criminals penetrate corporate networks to steal information. The personal data compromised include individuals' names, addresses, social security numbers, dates of birth, driver's licenses, passport numbers, and financial data. This information can then be used to commit crimes, including fraudulent unemployment claims (Goodin, 2008), fraudulent tax returns (McMillan, 2008), fraudulent loans (Hogan 2008), home equity fraud (Krebs, 2008), and payment card fraud. Consumers can also suffer the burden of increased loan interest rates, being denied utility services, civil suits or criminal investigation (Baum, 2004). While the consumer costs incurred from credit card fraud may be negligible, out of pocket expenses can reach thousands of dollars (FTC, 2007, Table 2).

As a result of these losses, in recent years U.S. policy makers have enacted laws that require organizations to notify individuals when personally identifiable information has been lost or stolen. As of late 2009, 45 states (as well as other countries around the world) have adopted data breach disclosure, or security breach notification, laws (Maurushat, 2009). Aside from two studies (one showing an improvement in firm practices (Samuelson, 2007), and another finding only a marginal reduction in consumer rates of identity theft (Romanosky et al., 2008)), however, the effects of data breach disclosure laws have yet to be rigorously studied.

One of the main intents of notification laws is to empower consumers to take action and mitigate their loss (Majoras, 2006). In addition, the possibility of loss from a breach and resulting costs from notification, it is argued, forces firms to internalize more of the cost of a data breach, thereby inducing them to increase their investment in security measures. This, in turn, is expected to reduce the probability, or magnitude, of future breaches. In short, data breach disclosure "drive[s] performance through transparency and oversight" (Mulligan, 2007).

However, critics argue that such laws inflict unnecessary costs for both firms and consumers if indeed firms already bear most of the loss (Rubin and Lenard, 2005) or when lost data is recovered before it is even accessed (Majoras, 2006). Moreover, when the risk of harm is low, unnecessary notification may desensitize individuals, preventing them from acting when a serious threat does exist (Majoras 2006). Further, consumers may be unable to properly respond to the breach notifications, as the notices may present a substantial cognitive and psychological barrier to tacking action, also causing them to under-react (Romanosky and Acquisti, 2009). Alternatively, news media and a burgeoning market of identity theft prevention services may breed panic and confusion, causing consumers to over-react by unnecessarily purchasing such products, increasing their expected costs.

But mandatory disclosure may also affect firms in conflicting ways. On the one hand, disclosure is costly. Firms will incur costs of notification, customer services operations (call centers, customer support), consumer redress (such as identity theft insurance or credit monitoring), legal fees, regulatory fines, and the potential loss of market valuation or lost business (customer churn) (GAO, 2007; Ponemon 2010). On the other hand, notifications may also cause consumers to take appropriate action and reduce their harm (either by preventing or mitigating identity theft) - this would lower the firm's own expected costs, because the amount of consumer harm that the firm internalizes is reduced.

In short, it is unclear whether disclosure would result in a net increase or decrease of firm, consumer, or overall social costs.

Using both analytical and numerical modeling, we show that even though firm costs will be higher under disclosure regimes, firms can be induced to increase their investment in care, which may lower social costs. Moreover, disclosure can induce consumers to increase their level of care, thus lowering their total costs. Finally, we find that the change in social costs are typically increasing in disclosure tax (costs imposed on the firm due to disclosure laws) and decreasing in consumer redress (compensation paid by the firm to the consumer). However, when the firm compensates consumers for only a small amount of loss, some disclosure tax may be necessary to optimally reduce social costs.

The next section discusses the literature related to information disclosure in IT security and the economics of (accident) law, which we leverage to frame information disclosure within the context of other common means of reducing externalities. We then define the costs involved in a data breach absent any legal regime, and illustrate how these costs change under mandatory breach disclosure. Next, we use analytical methods to determine the conditions under which disclosure reduces social costs. Finally, we provide discussion and empirical validation, followed by some model extensions and our conclusion.

2. Background

This research contributes to the information systems and information disclosure literatures as it relates to the economics of information security and the economics of accident law.

2.1 Economics of information security

The body of literature on IT security has been growing considerably in recent years, and much attention in this field has been paid to the disclosure of breaches, vulnerabilities, and software bugs. For example, Cavusoglu et al. (2008) examine a software vendor's incentive to distribute IT security patches. They compare a cost sharing policy where the vendor shares some burden of a firm's cost of applying software patches versus a liability policy where the vendor bears a portion of the cost of a firm's security incident as a result of an exploited IT vulnerability. Telang and Wattal (2007) find that a software vendors' stock price suffers when IT vulnerabilities are publicly disclosed. Both Telang et. al (2007) and Gandal et. al (forthcoming) provide a theoretical analysis of a vendor's incentive to disclosing IT vulnerabilities to consumers. Grossklags et al. (2008) provide a game-theoretical model describing a firm's incentive to reduce its losses either by increasing investment in security controls or by purchasing insurance.

Many have also empirically studied the effect of disclosing data breaches on stock market valuation. For instance, Campbell et al. (2003) find a negative effect on stock price for data breaches caused by "unauthorized access of confidential information." Cavusoglu et al. (2004) find that the disclosure of a security breach results in the loss of $2.1 of a firm's market valuation, on average. Acquisti et al. (2006) and Kannan (2007) both use event studies to measure the impact of a data breach on market price and while the former finds short-lived reduction of 0.6 percent on the day when the breach is disclosed, the latter finds no effect, on aggregate. Romanosky, Telang, Acquisti (2008) examine the effect of data breach disclosure laws on identity theft rates and find that the disclosure laws reduce identity theft by about 2%, on average.

Gordon et al. (2006) examine the effect of the Sarbanes-Oxley Act on the (attention paid to, and therefore the) disclosure of information security-related activities by firms. Similarly, Wang et al. (2009) find that firms that identify security threats in their 10-k filings using "action-oriented terms and phrases" are less likely to suffer a future data breach.

In general, the theoretical IS literature has focused on analysis of vulnerability disclosures even though the empirical literature has (just) started investigating the impact of breach notifications. However, no modeling analysis has specifically focused on mandatory disclosure of data breaches, which is our focus. As highlighted above, the trade-offs associated with disclosure

regimes are quite complex and nuanced for firms and consumers, making it not immediately discernible under what conditions such a regime may increase social welfare (lower social costs).

Empirical research has also been used to investigate the effect of mandatory disclosure polices on health outcomes (Jin and Leslie, 2003) and financial securities (Barth and Cordes, 1980). In particular, Beales et al. (1981) study information disclosure as a policy device and discuss the incentives for a firm to disclose product information under various market conditions. Specifically, they describe conditions that may lead a firm to over- or under-disclose either positive or negative product information. A great deal of research has also been devoted to disclosure (transparency) in IT software (end user license agreements; Good et al., 2005) and US policymaking (Fung et al., 2007).

Information disclosure, in general, can be used to increase market efficiency by improving consumer choice regarding products and product risk. When consumers act on this information firms are forced to respond and improve their product quality (or safety). This suggests that a policy of mandatory disclosure will be more useful when consumers either lack information or are misinformed, and that it will be less useful when consumers are either unable or unwilling to act (Beales et al, 1981). Moreover, disclosure will be more effective the more consumer action affects expected firm profits which may be especially true in more competitive environments. Fishman and Hagerty (2003) examine the incentives for firms to disclose product information and the differential effect on consumers who understand the notice, versus those who do not and that mandatory disclosure only benefits the informed consumer. We discuss the optimal and adverse effect of disclosure on consumer actions later in this article.

2.2 Economic analysis of accident law

From a modeling perspective, we analyze the economic impact of data breach disclosure regimes for firms, consumers, and society by leveraging the economics of accident law (Shavell, 1984; Landes and Posner, 1987; Kolstad, Ulen and Johnson, 1990).

Consider two cars on a roadway. Each driver engages in some level of care (prevention) and assumes some probability of an accident. Costs to the drivers include the actual cost of care, plus any expected damages as a result of an accident. Naturally, each driver will engage in a level of care that minimizes her private costs, which will be suboptimal when she does not bear the full cost of her actions. The objective of the social planner, therefore, is to devise a policy that induces drivers to take the socially optimal level of care, thereby minimizing total costs incurred by all parties.

For example, two common policy approaches are *ex ante* safety regulation and *ex post* liability. *Ex ante* safety regulation (i.e., a mandated standard) is meant to prevent accidents from occurring through the enforcement of minimum safety standards or operating (compliance) restrictions. An important characteristic is that sanctions can be imposed as soon as the regulations have been violated, even though no harm has yet occurred. For instance, drivers receive speeding tickets even without yet having caused an accident.[1] *Ex post* liability, on the other hand, is exercised after an accident has occurred. It is a legal device that enables victims to sue for damages, forcing injurers to internalize part of the harm they cause. Finally, a third approach, and the focus of this article, is information disclosure. Mandatory disclosure forces firms to reveal information about the risks of their products or services. As mentioned, the intent is to empower consumers to take action to mitigate potential harm, and to create a strong

[1] Other examples of mandated standards include industry operating licenses (driving, etc.), building safety and minimum health safety codes, fire proofing material, etc. Specific to consumer privacy, regulations include PCI (mandating minimum security controls on firms that process credit card transactions), HIPAA (mandating appropriate security controls on health care agencies and personal health information), Sarbanes-Oxley (section 404 requires, again, appropriate data protection measures for material financial information).

incentive for firms to improve their practices and avoid negative publicity or customer backlash.[2] Absent an information disclosure policy, only the firm can affect consumer harm, whereas under a mandatory disclosure policy, both parties can take action to reduce costs. Figure 1 illustrates these three approaches. The dashed vertical line represents an event that could lead to harm, such as a data breach, while the solid vertical line represents the actual harmful consequence, such as identity theft.

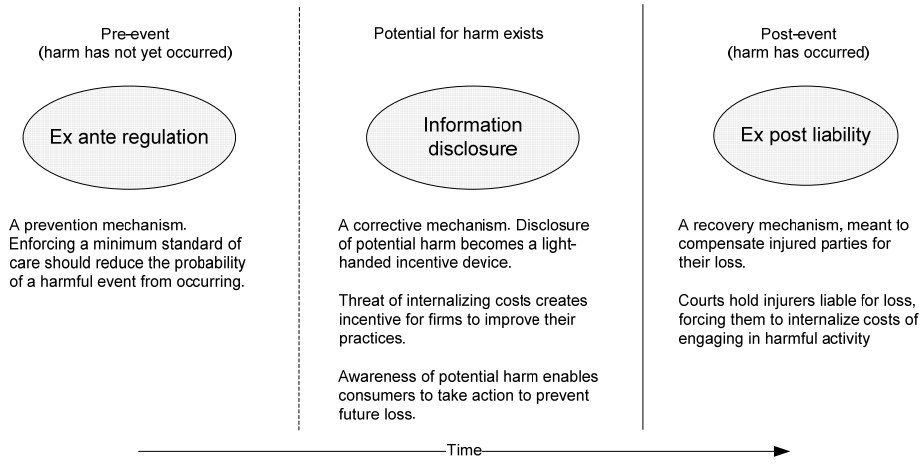

Figure 1: Three policy approaches[3]

Analytical (positive) analysis of these regimes often employ the economic analysis of law to investigate how policies and market structures drive incentives of economic agents. The positive analysis of tort law, for example, is concerned with the ways in which different liability rules induce one (unilateral) or two (bilateral) parties to take optimal precaution to avoid accidents (Shavell, 2004; Landes and Posner, 1987). This analysis shows that under unilateral-care accident models, both negligence and strict liability rules induce the optimal amount of care by the injurer. However, when care can be taken by both parties (bilateral-care accidents), these rules must be augmented with defenses of contributory or comparative negligence in order to also induce the victim to take optimal care.

Past research has directly compared *ex ante* safety regulation and *ex post* liability regimes to determine how they can either separately, or in combination, produce socially optimal outcomes. Shavell (1984) and Kolstad, Ulen and Johnson (1990) (hereafter KUJ) both demonstrate that social losses are minimized when regulation and liability regimes are employed together, but inefficient when used alone. Specifically, regulation is inefficient when the regulator lacks information about harm that occurs, or is uncertain about the appropriate minimum standard to set. Liability becomes inefficient when firms are not always held liable for the harm it creates, or when it is unable to pay for full damages (judgment proof).

KUJ discuss the means by which *ex ante* safety regulations can be used as either a substitute for, or complement with, liability. They show that the more uncertain is a firm's assessment of the probability of being held liable, the more likely it will be to under-invest in care. Their main conclusion is that under *ex post* liability it is inefficient to set the *ex ante* safety regulation at the

[2] Examples of existing disclosure policies include Toxic Release Information (mandates disclosure of hazardous material spills), FDA (requires that pharmaceutical companies notify the FDA in the event of harmful side effects of drugs), food nutrition labeling, fuel octane levels, cigarette warning labels, etc.

[3] Source: Romanosky and Acquisti (2009).

socially optimal level (where the marginal cost of prevention equals the marginal benefit of prevention).

Others have focused on modeling liability alone but incorporate actions by both the injurer and injured. Brown (1973) uses a non-cooperative game theoretic approach to model social loss when levels of prevention are endogenous to both parties under various liability conditions. For instance, the socially optimal level of prevention is achieved for most common-law negligence liability conditions, but not strict liability. Brown also finds an efficient level of investment for what he calls "relative liability" which enforces liability on the party that can reduce the harm at the lowest cost (often called the "least cost avoider"). Polinsky and Shavell (2006) admit that mandatory disclosure is better for consumers, but that disclosure in conjunction with a liability regime can lead to a suboptimal outcome because it reduces the firm's incentive to acquire information about product risks (through research and product testing).

In summary, much of the literature regarding information disclosure in information technology (IT) relates to the effect of data breaches on stock market valuation, and the incentives for firms to disclose IT vulnerabilities, yet does not address the social cost of data breach disclosure laws. The related literature on the economics of accident law, on the other hand, is vast in its investigation of the effects *ex ante* regulation and *ex post* liability, yet does not comparatively address information disclosure. Therefore, to our knowledge, this article is the first to theoretically analyze firm, consumer and social costs of a mandatory disclosure policy as it applies data breaches. We leverage the economic analysis of (accident) law to achieve a better understanding of how data breach disclosure laws drive incentives by firms and consumers to take more, less, or the socially optimal level of care.

Below we present the models which will provide the foundation for our comparison: firm, consumer and social costs with and without an information disclosure regime.

3. Economic Model

Our methodological approach follows the economic analysis of accident (tort) law, which is often concerned with developing policies to minimize accident costs. Therefore, we define cost functions of three parties: a firm (injurer), a consumer (victim) and the social planner, and we examine the conditions under which information disclosure induces behavior that may reduce social costs. We assume that the firm and consumer are rational economic agents, and that their objectives are to minimize their private costs by optimizing their level of care.

3.1 Basic model (no disclosure regime)

The basic firm and consumer cost model is presented below. First, the representative firm's amount of care, $x \geq 0$, represents the level of investment in all forms of security controls designed to reduce the probability or magnitude of a data breach. But care comes at a cost, defined as *c(x)*. These costs exist whether the breach occurs or not, and include technological (firewalls, encryption, ingress and egress filtering, authentication and authorization systems, and so forth) and administrative (user awareness, acceptable use policies, and so forth) investments.[4] We assume *c(x)* is increasing and convex in *x,* continuous and twice differentiable ($c' > 0$, $c'' > 0$, $c(0) = 0$, $c'(0) > 0$, $\lim_{x\to\infty} c(x) = \infty$).[5] We believe this to be a reasonable assumption, in that the same incremental level of care becomes more expensive as the level of care increases.

The probability of a breach, *p(x)*, reflects the probability that a breach occurs given a level of investment, x. We assume that *p(x)* is decreasing and convex in *x*, also continuous and twice

[4] We discuss the positive externalities of security investments later in this article.

[5] The condition that c'(0) = 0 eliminates the condition where a firm would have no incentive to invest in securing its data assets against breaches. Practically, it represents the case when some security measures may be applied for zero cost, such as changing default password settings on enterprise applications.

differentiable ($p' < 0, p'' > 0, p(0) = 1, \lim_{x\to\infty} p(x) = 0$). Following the economic analysis of tort law, our model focus on a representative firm, which means that we do not differentiate between firms and we do not consider how one firm may be more attractive than another to cyber-criminals (which may raise its probability of a breach).

Firms suffer a cost of investigation, $i > 0$. This is the cost incurred by the firm to investigate the cause and scope of a data breach. Regardless of the threat or requirement to disclose a breach, firms must still respond to security incidents, determine their cause, repair damaged IT systems and ensure business services are fully operational (Lemos, 2009).

Next, we address consumer costs. First, note that not all breaches will result in identity theft. For instance, backup tapes of financial or medical information may be lost, or computer hardware may simply be stolen to be resold as parts. Therefore, we only consider consumer harm caused by data breaches.[6] We assume, in this basic model with no disclosure, that because consumers are not informed about the potential for identity theft, they are unable to take action to prevent or mitigate any loss. (Since this is a key feature of disclosure laws, consumer care is modeled in the next section.) We consider that the total consumer cost of identity theft is generally comprised of two components. First, there is the actual loss represented by the amount of money stolen by an attacker. Estimates of these values range from $0 (for example, in credit card fraud where one's bank may fully reimburse the victim for any money stolen), to thousands of dollars or more (FTC, 2007, Table 2). Second, there are costs incurred directly by the victim, such as loss or denial of financial or utility (telephone, electrical, etc.) services, time and effort required to recover and restore one's credit, higher interest rates, being subject to a civil suit or even criminal investigation (Baum, 2004, Table 3). We define these costs as $h \geq 0$.[7]

We can now define the firm's objective function, and the consumer and social cost functions absent any legal mechanism. The firm's objective is to determine the level of care, x, that minimizes its total costs:

$$F(x) = c(x) + p(x)\,i \tag{1}$$

Since consumer harm is exogenous to the consumer, their loss function is

$$C(x) = p(x)h \tag{2}$$

The social loss is, therefore, given by

$$S(x) = c(x) + p(x)\left[i + h\right] \tag{3}$$

Below, we investigate the changes in firm and consumer behavior under a policy of mandatory information disclosure.

3.2 Data breach disclosure model

Mandatory breach disclosure implies that a firm that has lost its consumers' data (or whose consumers' data has been stolen, or otherwise compromised) needs to notify the latter. One

[6] We discuss the data relating to the portion of identity theft due to data breaches in Section 6.

[7] Strictly speaking, the theft of money represents a cost only to the victim, but not a social cost. The reason is that it represents, in essence, only a transfer of money between the victim and criminal. Therefore the component of social cost due to consumer loss would reflect the victim's psychological harm, time and effort, increased interest rates, etc. For simplicity, we combine these effects into a single variable h, recognizing that separating them would increase complexity without revealing additional insights.

consequence of information disclosure is to transform unilateral-care accidents into bilateral-care accidents. Therefore, under disclosure, consumers are enabled to take action to reduce their harm. For example, once notified, consumers can closely monitor their credit reports for signs of fraudulent activity; they can stop transactions or cancel financial or retail accounts that they believe to be compromised; they can notify their banks and place credit freezes or fraud alerts on their credit reports, purchase identity theft insurance, or file individual or class action lawsuits[8] in order to recover actual or potential losses from identity theft. We consider that any or all of these actions will reduce their expected loss from identity theft, and we call the level of consumer care $y \geq 0$.

However, there is effort (cost) of care. Total costs include both the actual consumer loss from identity theft as well as psychological costs of understanding, interpreting, and deciding how to react with the information provided by the notification, financial costs, and time spent addressing potential harm (Romanosky and Acquisti, 2009).

We assume that the marginal benefit received from taking action is decreasing in level of care, while the marginal cost is increasing (we later consider how consumers may under- or over-react to notification). Together, these behaviors produce the convex consumer harm function *h(y)*, shown in Figure 2.

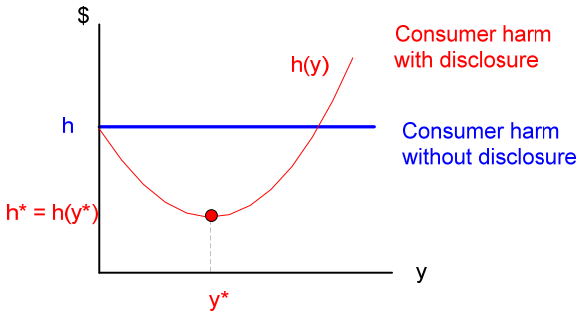

Figure 2: Consumer harm as a function of consumer care

Note that consumer harm is bounded on the left at $h(y=0) = h$, illustrating how, absent any action, consumer harm without disclosure is equal to consumer harm with disclosure. Next, this convex loss function naturally exhibits a minimum, y^*, which represents the smallest consumer harm possible under disclosure, $h_D = h^* = h(y^*) < h$. Finally, we assume that the increasing cost of care will eventually dominate any benefits, resulting in consumer loss under disclosure that will equal, then surpass, consumer loss without disclosure, $\lim_{y \to \infty} h(y) = \infty$.

We now turn to firm costs under mandatory disclosure, making two important distinctions: i) we separate costs that are dependent and independent of consumer action (or consumer harm), and ii) we distinguish between costs that increase social loss (e.g., deadweight loss) versus those that represent a transfer of funds between the firm and the consumer.

First, as before (absent a disclosure regime) a firm will incur the cost of investigating a data breach, repairing any IT systems, and restoring business services, $i > 0$. We assume that this parameter is independent of a disclosure regime (i.e., firms must identify the cause of the breach regardless of a disclosure requirement).

Second, firms will incur many costs because of the disclosure policy itself. These include the costs of legal fees related to determining whether to disclose the breach or not, providing customer support (call center) to respond to inquiries, litigation holds (document preservation),

[8] Private and class action lawsuits have been filed in response to many breaches, including Ameritrade (Kravets, 2008), Starbucks (McMillian, 2009), or Heartland (McGlasson, 2009). TJX paid $525,000 to banks (Kaplan, 2009) and allocated $256M to cover costs of its breach (Kerber, 2007).

forensic investigation, and human resources (employee termination or disciplining) (Ponemon, 2010, Table 3, p21).[9] Also included are costs of customer notification, public relations campaigns, and regulatory fines or fees imposed by state or federal regulatory agencies (FTC, State health agencies). In cases of breaches of payment cards, VISA and MasterCard also fine merchants (indirectly, through their acquiring bank), and banks are able to increase a merchant's interchange fees (creating, essentially, a tax imposed by the bank on the merchant per transaction).[10] Regulatory sanctions sometimes also require increased security precautions and auditing (often up to 20 years afterward) and can impose fines for consumer redress (to compensate consumers for any losses related to the breach). In addition, the sunlight effect (Romanosky et al., 2008) implies that either consumers or the market (buyers and sellers of a firm's stock) can punish firms for their bad practices, destroying market valuation or cause loss of business and revenue.[11] Indeed, one survey claims that consumers sever relationships as a result of a data breach and that the cost of lost business can reflect up to 40% of the total cost of a data breach (Ponemon, 2010, Table 3, p21). All of these costs are borne strictly by the firm and represent a social (deadweight) loss and what we will consider collectively as the "disclosure tax," $d \geq 0$.[12]

Another consequence of the data breach disclosure policy (intentionally or not) is to force firms to bear some portion of consumer loss. For example, firms are often encouraged or required to provide consumer redress through regulatory fines (e.g., FTC v. Choicepoint, 2005) or free credit monitoring services. Moreover, consumers continue to file individual and class action lawsuits in the hopes of recovering losses from data breaches,[13] and in this regard, the firm would become liable for consumer losses. The extent to which a firm should bear more or less consumer harm is currently under great debate, however. On one hand, privacy advocates consider it a great failure of the justice system when claims filed against firms for data breaches are promptly dismissed. On the other hand, the governor of California recently vetoed a data breach bill claiming that firms already bore enough liability, "the marketplace has already assigned responsibilities and liabilities that provide for the protection of consumers" (Schwarzenegger, 2007). Moreover, the legal issue of sufficient liability is still unresolved. At the time of this writing, the Maine Supreme Court is addressing whether time and effort mitigating potential identity theft is compensable under state tort law (Zetter, 2009). And so, we represent the portion of consumer harm borne by the firm as $\lambda h(y)$, with $0 \leq \lambda \leq 1$. A value of $\lambda = 1$ implies that the firm fully compensates the consumer for their loss, while a value of $\lambda = 0$ implies that the firm bears no consumer loss. The amount of consumer harm borne by the consumer therefore becomes $[1-\lambda]\, h(y)$. That the firm bears a portion of consumer harm (which is dependent on the level of care by the consumer) is a reflection of compensatory (i.e., not punitive) damages under US tort law whose purpose is to reimburse victims for the loss they would otherwise not have suffered but-for the accident. That is, damages are awarded such that the victim is made 'whole' (Buckley

[9] Sources include conversations between an author of this paper and US lawyers who specialize in data breach litigation.

[10] E.g., Heartland paid over $6 million to payment card processors as a result of its breach (Goodin, 2009). Moreover, firms may also face increased costs per transaction, called interchange fees. For example, Kaiser was fined twice by state health organizations for $187,500 (Zick, 2009).

[11] Certainly the ability for a consumer to change firms is not uniform. Shopping at a new grocery or retail store, or gas station, for instance, may be simple, while changing financial institutions (mortgage, banking, credit card), schools, utility companies or employer would incur much higher switching costs (perhaps prohibitive).

[12] While we have combined all firm costs into a single variable, d and labeled it a disclosure "tax," we recognize that it is not strictly a tax as would be imposed by a government authority.

[13] See Romanosky and Acquisti (2009) for a discussion of the different ways in which consumers bring action against firms for data breaches, but are often unsuccessful.

and Okrant, 2003).[14] An important property of this cost is that it represents, in essence, a transfer payment from the firm to the consumer (i.e., consumer costs paid by the firm to the consumer) that lowers total consumer loss, but does not change social cost.[15]

Finally, we consider that these costs are a function of the size of a breach, though for simplicity we normalize all costs as a loss per record (or per account). That is, *d* represents the cost to the firm per account lost or stolen, and *h* represents consumer loss per individual from a data breach.

We now define firm and consumer objective functions and a social cost function under mandatory disclosure. The firm's objective, again, is to determine the level of care, *x*, that minimizes its total costs

$$F_D(x) = c(x) + p(x)\left[i + d + \lambda\, h(y)\right] \tag{4}$$

With consumer harm now endogenous, the consumer's objective is to also determine their level of care, *y*, that minimizes total costs

$$C_D(x, y) = p(x)\left[1 - \lambda\right]h(y) \tag{5}$$

The social cost function is therefore,

$$S_D(x, y) = c(x) + p(x)\left[i + d + h(y)\right] \tag{6}$$

The firm's cost function is optimized at $\tilde{x}$ without information disclosure and at $\tilde{x}_D$ with information disclosure. Similarly, the social cost function is optimized at x^* without disclosure and x_D^* with disclosure. A summary of variables is show in Table 1.

Table 1: Variable descriptions

Variable	Description
x, y	Level of care (firm, consumer)
c(x)	Cost of firm care
p(x)	Probability of a data breach
$\tilde{x}, \tilde{y}$	Privately optimal level of care without disclosure (firm, consumer)
$\tilde{x}_D, \tilde{y}_D$	Privately optimal level of care with disclosure (firm, consumer)
x_D^*, y_D^*	Socially optimal level of care with disclosure (firm, consumer)
x^*	Socially optimal level of care without disclosure
h	Consumer harm without disclosure
h_D	Consumer harm with disclosure
*h**	Optimal consumer harm with disclosure
i	Cost to the firm of investigating a breach

[14] Another example is provided in the ruling of Spangler v. Helm's New York-Pittsburgh Motor Express, 396 Pa. 482, 485, 153 A.2d 490, 492 (1959), "As between the innocent victim of a wrong and the person who accomplished the wrong, the law imposes on the malfeasor the obligation to make the victim whole in every phase in which the victim has suffered…"

[15] For example, Choicepoint paid $10 million in civil penalties and $5 million in consumer redress (Brodkin, 2007), while the Veterans Affairs agency agreed to pay $20 million in consumer redress, including credit card monitoring in response to a breach (Pulliam, 2007).

d	Disclosure tax
λ	Portion of consumer harm borne by the firm
F, C, S	Total cost without disclosure (firm, consumer, social)
F_D, C_D, S_D	Total cost with disclosure (firm, consumer, social)

The cost equations with and without disclosure are presented in Table 2 for convenience.

Table 2: Cost equations

Cost	No Disclosure	Mandatory Disclosure
Firm	Eq. (1) $F(x) = c(x) + p(x)\,i$	Eq.(4) $F_D(x) = c(x) + p(x)\left[i + d + \lambda\, h(y)\right]$
Consumer	Eq.(2) $C(x) = p(x)h$	Eq.(5) $C_D(x, y) = p(x)\left[1 - \lambda\right]h(y)$
Social	Eq.(3) $S(x) = c(x) + p(x)\left[i + h\right]$	Eq.(6) $S_D(x, y) = c(x) + p(x)\left[i + d + h(y)\right]$

4. Effect of Disclosure on Consumer and Firm

Below, we analyze the effect of a mandatory disclosure policy first on the consumer and then on the firm.

As shown previously in Figure 2, a rational consumer will engage in a level of care that minimizes their private costs. That is, the consumer will take action y until $\tilde{y}_D = y^*$ (until their privately optimal level of care under disclosure equals the socially optimal level of care), resulting in loss of $h_D = h(y^*) = h^*$. Since consumer action without disclosure is assumed to be zero, it is trivial to conclude that disclosure increases consumer care. Moreover, total consumer loss under disclosure will always be lower under the current assumptions, because not only will the consumer incur less harm through their actions, but they also now bear only a fraction of that harm.

Proposition 1a*: Under a policy of mandatory disclosure, consumers will take more care, but their costs will be strictly lower.* That is, $\tilde{y} < \tilde{y}_D = y^*$, and $C_D(x, y^*) < C(x)\ \forall\, x$.

Proof: Since consumer care without disclosure is assumed to be zero, and consumer harm, *h(y)* is decreasing in y until y^*, consumer care must be greater under disclosure. Since, $h(y^*) < h$ and $0 \le \lambda \le 1$, it is obvious that consumer costs with disclosure must be lower than consumer costs with disclosure for all *p(x)*.

Next, we compare the firm's optimal level of care with the socially optimal level of care. By construction, the firm will not internalize the full amount of its actions, and therefore under-invest in care with or without a mandatory disclosure policy. Further, when disclosure becomes costly to the firm, either through the many costs associated with notification, *d*, or when the firm bears some portion (however small) of consumer harm, the firm will take more care in order to minimize the total costs of a data breach. Figure 3 illustrates how the marginal cost of care remains constant, but greater potential costs due to a breach (due to *d* and $\lambda\, h(y)$) will increase the marginal benefit (of loss avoided).

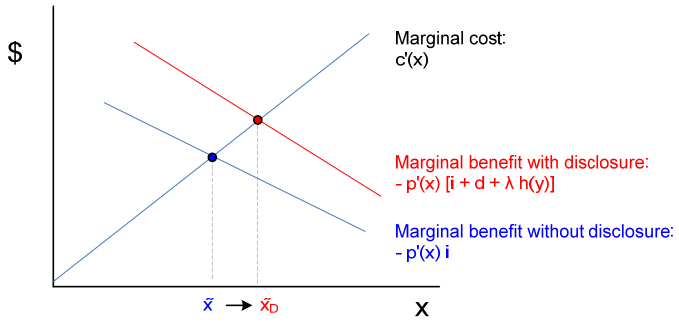

Figure 3: Increasing firm costs under disclosure

Finally, as a result of these additional costs, and the increased level of care, firm costs under a disclosure policy will always be higher, relative to no disclosure as illustrated in Figure 4. The left panel shows how firm costs are greater for any level of care, while the right panel shows how firm costs are greater especially for the cost minimizing level of care.

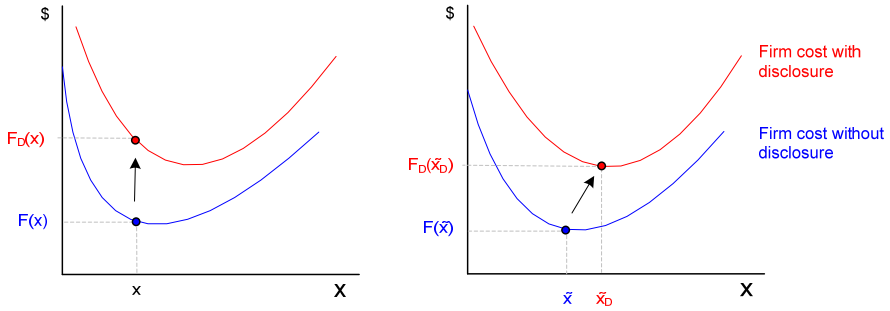

Figure 4: Increasing firm costs

We summarize the three firm propositions below.

Proposition 1b*: A firm will under-invest in care either with or without a policy of mandatory disclosure (relative to the socially optimal level of care). That is,* $\tilde{x} < x^*$*, and* $\tilde{x}_D < x_D^*$.

Proof: The optimal firm investment is $\tilde{x}$, therefore, $\tilde{x}$ satisfies $c'(x) + p'(x)i = 0$. That there is a solution to this and similar equations is a consequence of the condition that $c'(0) = 0$. The social cost at $\tilde{x}$ is decreasing: $S'(\tilde{x}) = c'(\tilde{x}) + p'(\tilde{x})[i + h] = p'(\tilde{x})h < 0$. Since the social cost is convex and decreasing at $\tilde{x}$ it must be that $\tilde{x} < x^*$ where x^* is the socially optimal level of care. Similarly, $\tilde{x}_D < x_D^*$. Eq. (4) shows how the firm will only invest in the socially optimal amount of care under disclosure when $\lambda = 1$.

Proposition 1c*: The firm will invest more in care when forced to disclose a data breach, relative to a non-disclosure regime. That is,* $\tilde{x}_D > \tilde{x}$.

Proof: $\tilde{x}$ is located at the unique minimum of the firm's convex cost function, which is determined by solving $c'(\tilde{x}) + p'(\tilde{x})i = 0$. Likewise, $\tilde{x}_D$ is found by solving $c'(\tilde{x}_D) + p'(\tilde{x}_D)[i + d + \lambda h(y)] = 0$. Since $d > 0$ and all other parameters are non-negative, we can write,

$$-\frac{c'(\tilde{x})}{p'(\tilde{x})} = i \;<\; i + d + \lambda h(y) = -\frac{c'(\tilde{x}_D)}{p'(\tilde{x}_D)} \tag{7}$$

Which gives,

$$-\frac{c'(\tilde{x})}{p'(\tilde{x})} \;<\; -\frac{c'(\tilde{x}_D)}{p'(\tilde{x}_D)} \tag{8}$$

Since *p(x)* and *c(x)* are convex, *c'(x)* is increasing while *-p'(x)* is decreasing in *x* which implies that the function - *c'(x)* / *p'(x)* is increasing in *x*. Therefore, it must be that $\tilde{x}_D > \tilde{x}$.

Proposition 1d*: Firm costs will be higher under disclosure: i) for any given level of care, and ii) at the firm's optimal (cost minimizing) level of care. That is,* $F_D(x) > F(x)$, and $F_D(\tilde{x}_D) > F(\tilde{x})$.

Proof: Because $\tilde{x}$ minimizes $F(x)$, $F(\tilde{x}) < F(\tilde{x}_D)$, and because $F(x) < F_D(\tilde{x}) \;\forall\; x$, $F(\tilde{x}_D) < F_D(\tilde{x}_D)$. Therefore $F(\tilde{x}) < F_D(\tilde{x}_D)$.

5. Effect of Disclosure on Social Costs

Now that we have determined the effects of mandatory disclosure on consumer and firm behavior, we turn our attention to our primary research interest: investigating the conditions under which a data breach disclosure regime reduces social cost, relative to no disclosure, and when this reduction is optimized. Since we are interested in the *change* in social cost due to the disclosure policy, our quantity of interest is

$$\Delta S = S_D(x_D) - S(x). \tag{9}$$

When the sign of Eq. (9) is negative, we conclude that information disclosure reduces social cost. That is, the more negative is ΔS, the greater is the reduction (i.e., improvement) in social cost due to mandatory information disclosure. In order to properly evaluate this expression, however, we must be sure to assess it at the firm's cost minimizing level of care under each policy. That is

$$\Delta S = S_D(\tilde{x}_D) - S(\tilde{x}).^{16} \tag{10}$$

The reason is that when a social planner decides to implement a disclosure policy, the firm will react by adjusting its level of care based on its new (private) cost structure. Given this new level of care, data breaches will occur with some new probability. Given their expected harm, consumers then react by taking optimal care. We solve this interaction by backward induction, representing it as a two-stage game where the consumer acts only after being notified. Given that *p(x)* is now exogenous to their strategy, the consumer perceives *p(x| data breach)* = 1. Hence, we assume that the consumer takes optimal care, $y_D = y^*$, which produces harm $h^*=h(y^*)$ (as previously illustrated in Figure 2). We then solve for the firm's optimal care, given h^*. Finally, we evaluate Eq. (10), the change in social costs, evaluated *at the firm's optimal level of care*.

[16] Read as: the change in social cost is equal to the social cost with disclosure (evaluated at the firm's cost-minimizing level of care with disclosure), minus the social cost without disclosure (evaluated at the firm's cost-minimizing level of care).

5.1 Movement of social cost curve vs. movement along the curve

Now that we have defined our quantity of interest, we must discuss two different behaviors that affect the social cost curve, and therefore ΔS. The first is the movement *of* the social cost curve and the second is the movement *along* the curve.

The movement of the social cost curve is affected by each of the parameters in the cost functions, *d, h*, and *h**,and *i*. As any of these parameters increases, the social cost curve will rise vertically because the social cost for any given value of *x* is greater (note the shift at the vertical intercept), as shown in the left panel of Figure 5.

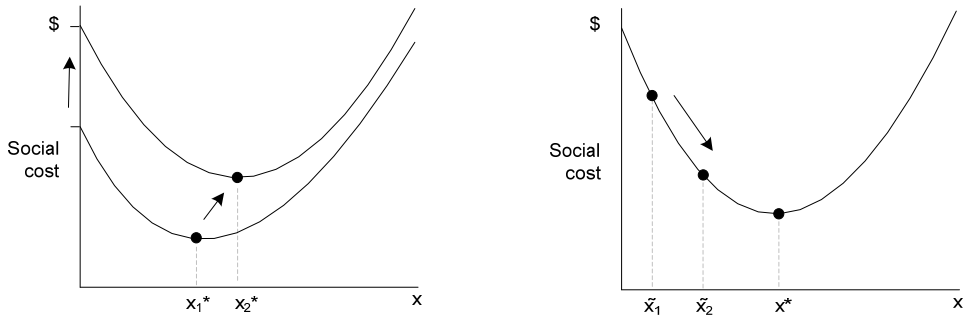

Figure 5: Movement of, and along, the social cost curve

Moreover, an increase in these parameters also increases the marginal benefit of harm avoided, driving the socially optimal level of care to the right ($x_2^* > x_1^*$). The net result is that as the parameters increase, the social cost curve is driven upwards, and to the right. We also notice how the difference between the curves is maximal at the vertical intercept (where the difference is equal to the change in the parameter values) and is decreasing in *x* (where the limit of the difference equals zero as *x* approaches infinity).

Next, the movement along the curve is determined by the portion of the total social cost borne by the firm (relative to the consumer). Holding the position of the social cost curve constant, the movement along the curve can be driven entirely by λ. A change in λ does not affect the position or shape of the social cost function, but simply changes the firm's optimal level of care. Specifically, when λ increases, the location of the firm's level of care "slides down" the social cost curve, approaching the socially optimal level of care as shown in the right panel of Figure 5 ($\tilde{x}_2 > \tilde{x}_1$). This occurs because the change in λ represents a transfer of cost between the firm and consumer. As mentioned, when λ = 0, the consumer bears all the cost (the firm bears none). On the other hand, when λ = 1, the firm bears the full cost (eliminating the externality) causing the firm's cost minimizing level of care to equal the socially optimal level of care.

We described how the movement of the social cost curve is driven by the parameters in the social cost functions. However, given our interest in the difference between the social costs with and without disclosure, we next focus our attention on the difference between these parameters as shown by expanding Eq. (10) as,

$$\Delta S = c(\tilde{x}_D) + p(\tilde{x}_D)[i + d + h^*] - \left(c(\tilde{x}) + p(\tilde{x})[i + h]\right) \qquad (11)$$

That is, we are interested in the difference between *d* + *h** (the sum of the disclosure tax and consumer harm under disclosure) and *h* (consumer harm absent disclosure). (We assume that i, the cost of a breach investigation is unchanged by a disclosure policy.) However, observe that *h* – *h** represents the *change in consumer harm*, and is clearly an important outcome measure of the

policy intervention. For instance, when $h - h^* > 0$, we conclude that consumer harm has, indeed, declined due to consumer action. And so we can rewrite Eq (11) as,

$$\Delta S = \Delta\tilde{c} + \tilde{p}_D(d - \Delta h) + \Delta\tilde{p}[i + h] \qquad (12)$$

Where $\Delta\tilde{c} = c(\tilde{x}_D) - c(\tilde{x})$, $\Delta\tilde{p} = p(\tilde{x}_D) - p(\tilde{x})$, $\tilde{p}_D = p(\tilde{x}_D)$, *and* $\Delta h = h - h^*$

Therefore, rather than analyzing the effect of *d, h, h** separately (plus λ) on the change in social cost, we can instead simplify the comparison by holding Δ*h* fixed and evaluating the change in social cost over a range of *d*, with special attention to $d < \Delta h$, $d = \Delta h$, *and* $d > \Delta h$. Intuitively, this highlights the comparison between the private *cost* imposed on a firm from a disclosure policy and the social *benefit* achieved from reduced consumer harm.

Below, we explore the change in social cost as a function of *d* and λ and therefore examine $\Delta S(d, \lambda)$ for $d < \Delta h$, $d = \Delta h$, *and* $d > \Delta h$, and λ [0,1]. Note that λ does not appear explicitly in Eq. (12) but is implicit in $c(\tilde{x}_D)$ and $p(\tilde{x}_D)$.

5.2 Illustration of ΔS(d,λ)

We begin the analysis by illustrating the behavior of $\Delta S(d, \lambda)$ over *d* (that is, different levels of disclosure tax) for values of λ equal to 0, 0.1 and 1 (that is, different portions of consumer harm borne by the firm), as shown in Figure 6.[17] We assume in this section that disclosure has made the consumer no worse off, implying that $h_D = h^* < h$. We consider the consequence of $h_D > h$ later in this article.

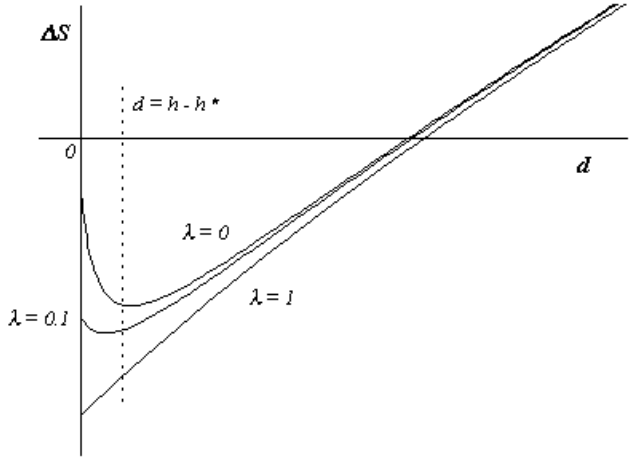

Figure 6: Graphical example of ΔS(d, λ)

Observe that the optimal level of disclosure occurs when the firm bears the full cost of consumer harm (lowest curve, $\lambda = 1$), and that the first-best social cost (optimal disclosure) occurs for the trivial case along this line (when the firm bears all consumer cost) and when the disclosure tax is zero, $\Delta S(d=0, \lambda=1)$.

Moreover, while ΔS is clearly negative for all values of $d \leq h - h^*$, even when the disclosure tax is greater than the consumer benefit ($d > h - h^*$), disclosure can still reduce the social cost

[17] For illustrative purposes we choose the following simple cost and probability functions: $c(x) = x^2$ and $p(x) = 1/(1 + x)$.

(the curves pass through the rectangle bounded on the left by the line $d = h - h^*$ and on the top by the horizontal axis).

Finally, observe that when the firm internalizes only a very small portion of consumer harm (shown by the curves $\lambda=0$, and $\lambda=0.1$), social cost is decreasing in d, implying that some disclosure tax is necessary to achieve the greatest reduction in social cost.

To begin the analysis, we can specify several basic properties of ΔS which are formalized in the Appendix: it can be proved that ΔS is well defined and continuous for all d and λ and that all values of ΔS will lie between $\Delta S(d,0)$ and $\Delta S(d,1)$. Moreover, ΔS is increasing without bound in d and $\Delta S(d,0)$ and $\Delta S(d,1)$ converge as d tends to infinity.

5.3 When is disclosure optimal?

5.3.1 First-best social cost

We first evaluate the case when disclosure produces the greatest reduction in social cost and therefore the first-best case. As shown in Figure 6, the largest negative value of ΔS, the optimal level of disclosure, is achieved when the disclosure tax is zero and the firm bears the full amount of consumer loss. When the firm bears the full extent of consumer harm its cost-minimizing level of care equals the socially optimal level of care ($\tilde{x}_D = x_D^*$). Moreover, with no disclosure tax, the change in social cost is maximized because the social cost under disclosure is minimized. This is a trivial outcome in that it describes a condition of zero externality and no deadweight loss caused by a tax.

Proposition 2*: The first-best social cost occurs when the firm bears the full consumer loss and there is no disclosure tax. That is,* $\Delta S(0,1)$.

Proof: By Lemma 1 (see the Appendix for accompanying Lemmas), we know that for any given d, $\Delta S(d,1) < \Delta S(d,\lambda)\ \forall\ \lambda \in [0,1)$. Therefore, we need only consider the case of $\lambda = 1$, and we observe that $\Delta S(d,1)$ is strictly increasing in d,

$$\frac{d\Delta S}{dd} = \frac{d\Delta \tilde{c}}{dd} + \frac{d\Delta \tilde{p}}{dd}(i+h) + \frac{d\tilde{p}_D}{dd}(d - \Delta h) + \tilde{p}_D \tag{13}$$

But we have,

$$\frac{d\tilde{c}_D}{dd} = \frac{\partial \tilde{c}_D}{\partial x}\frac{\partial \tilde{x}_D}{\partial d};\ \frac{d\tilde{p}_D}{dd} = \frac{\partial \tilde{p}_D}{\partial x}\frac{\partial \tilde{x}_D}{\partial d} \tag{14}$$

Which gives,

$$\frac{d\Delta S}{dd} = \left(\frac{\partial \tilde{c}_D}{\partial x} + \frac{\partial \tilde{p}_D}{\partial x}(i+d+h^*)\right)\frac{\partial \tilde{x}_D}{\partial d} + \tilde{p}_D \tag{15}$$

$$\frac{d\Delta S}{dd} = \left(-\frac{\partial \tilde{p}_D}{\partial x}(i+d+\lambda h^*) + \frac{\partial \tilde{p}_D}{\partial x}(i+d+h^*)\right)\frac{\partial \tilde{x}_D}{\partial d} + \tilde{p}_D \tag{16}$$

$$\frac{d\Delta S}{dd} = \frac{\partial \tilde{p}_D}{\partial x}\frac{\partial \tilde{x}_D}{\partial d}(1-\lambda)h^* + \tilde{p}_D \tag{17}$$

And for $\lambda = 1$ we have

$$\frac{d\Delta S}{dd} = \tilde{p}_D > 0 \qquad (18)$$

Thus, $\min_{d,\lambda} \Delta S(d,\lambda) = \min_{d} \Delta S(d,1) = \Delta S(0,1)$.

Clearly, this case of optimal disclosure would practically only be achieved by a benevolent dictator, able to force a firm to bear all consumer harm while making disclosure costless to the firm. The implication is that a policy maker (or this benevolent dictator) would not need to further punish the firm with a disclosure tax since the firm is already bearing the full extent of consumer harm.

Next, we consider the means by which the minimum social cost can be achieved when the globally optimal conditions are not met, by investigating the interaction between the disclosure tax, *d*, and portion of consumer harm borne by the firm, λ.

5.3.2 Second-best social cost

As just shown, ΔS becomes more negative (i.e., the disclosure policy is improving), first as λ increases to 1, then as *d* decreases, reaching a minimum (the first-best social cost) at $\lambda = 1$ and $d = 0$. However, as shown in Figure 6, ΔS is minimized at a positive value of *d* when λ is small. In other words, a disclosure tax of zero (costless disclosure) was only optimal when the firm internalized more than a certain threshold of consumer costs. Therefore, we conclude that there may exist a threshold, λ_T, for which social costs are increasing in disclosure tax, *d*, when $\lambda > \lambda_T$, and below which social costs are decreasing in *d*, when $\lambda < \lambda_T$. Furthermore, when optimal consumer harm under disclosure, *h**, is sufficiently large, $\lambda_T > 0$, implying some disclosure tax is necessary to minimize social cost.

Proposition 3: *There may exist a threshold, λ_T, such that some disclosure tax is necessary in order to minimize social cost. Moreover, when h* is large enough, there will exist a λ_T such that* $d\Delta S(0,\lambda)/dd < 0$ for $\lambda < \lambda_T$.

Proof: We have already shown that $d\Delta S(0,1)/dd > 0$. Next, we can show that $d\Delta S(0,0)/dd$ may be less than zero,

$$\left.\frac{d\Delta S}{dd}\right|_{d=0,\lambda=0} = \frac{\partial \tilde{x}_D}{\partial d}\frac{\partial \tilde{p}_D}{\partial x}h^* + \tilde{p}_D \qquad (19)$$

The first term on the right hand side of Eq. (19) is negative while the second is positive, therefore their sum may be less than zero. If $d\Delta S(0,0)/dd < 0$, then there exists $0 < \lambda_T < 1$ such that $d\Delta S(0,\lambda_T)/dd = 0$ since $d\Delta S/dd$ is a continuous function of λ.

Next, observe that the sum of the terms on the right hand side of Eq. (19) will be less than zero if *h** is large enough. Because $d = 0$ and $\lambda = 0$, $\tilde{x}_D$ does not depend on *h**, and neither does the product $\frac{\partial \tilde{x}_D}{\partial d}\frac{\partial \tilde{p}_D}{\partial x}$, and so the negative term in Eq. (19) will dominate the positive term as *h** increases. The expression for the threshold value of λ_T is then,

$$\lambda_T = 1 - \frac{\tilde{p}_D}{\frac{\partial \tilde{x}_D}{\partial d}\frac{\partial \tilde{p}_D}{\partial x}h^*} \tag{20}$$

The practical implication of this result, and one of the key findings of this article, is that when firms bear a small enough portion of consumer loss, some disclosure tax may be necessary in order for social costs to reach a minimum. However, the policy maker would not need to punish the firm with a disclosure tax, if the consumer can very effectively protect themselves (demonstrated by a low h^*). Note that this doesn't imply that the disclosure policy should not be adopted, because it is the policy, itself, that enables consumer notification allowing consumers to take care. Further, the disclosure policy, as we've shown, forces the firm to increase its level of care, which may also be a useful policy objective.

We now focus our attention on the difference between the firm's costs of disclosure tax and the benefit of reduced consumer harm from an information disclosure policy. This comparison is relevant because, practically speaking, a policy maker may not have the ability to force the firm to bear all consumer losses (due to legal and market constraints) and may likely only have limited control of the firm's increased costs under a disclosure policy. However, the policy maker may indeed face the decision of whether or not to pass a data breach disclosure law, therefore creating costs d and λ for the firm. In the next sections we examine the conditions under which a policy maker should and should not introduce disclosure legislation.

5.4 When is disclosure preferred?

5.4.1 Disclosure tax is less than or equal to consumer benefit

Let us first evaluate the change in social cost when the disclosure tax is less than or equal to the benefit from lower consumer harm – that is, when $d \leq \Delta h$. First, when $d < \Delta h$ the social cost curve with disclosure will always lie below the social cost curve without disclosure as shown in Figure 7.

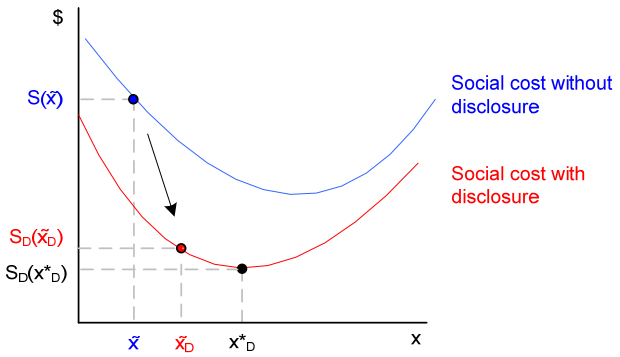

Figure 7: Decreasing social costs when d < Δh

This is a simple case: the social cost with disclosure is less than the social cost curve without disclosure for every x. Moreover, by Proposition 1c, the firm's level of care under disclosure is greater than the level of care without disclosure; therefore, the social cost must be decreasing.

Next, when the disclosure tax is equal to the reduction in consumer harm *($d = \Delta h$)*, the social cost functions with and without disclosure become equivalent. That is, Eqs. (3) and (6) become identical, and the social cost for any given level of care is the same both with and without disclosure. Moreover, as given by Proposition 1c, the firm will increase its private (optimal) level

of care, which implies that the position on the social cost curve will tend downward, toward the socially optimal level of care, thus reducing the social cost, as shown in Figure 8.

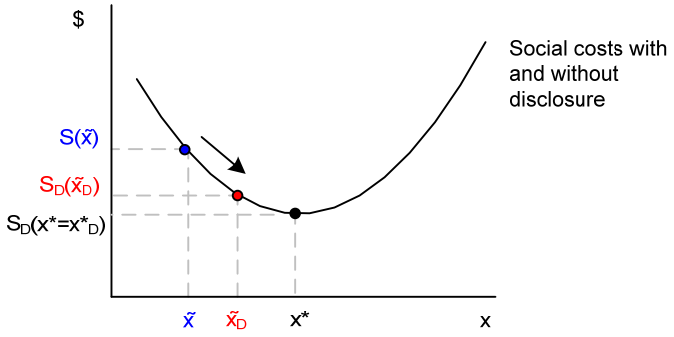

Figure 8: Decreasing social cost when *d* = Δh

Proposition 4*: When the disclosure tax is less than or equal to the reduction in consumer harm, social cost will always be lower. That is,* $S_D(x) \leq S(x)\ \forall\ (d,\lambda) \in\ ]0,\infty\ \times\ ]0,1]$,

Proof: As previously shown $\tilde{x} < x^*$, $\tilde{x}_D < x_D^*$ and $\tilde{x} \leq \tilde{x}_D$, therefore, $x^* = x_D^*$, and $\tilde{x} \leq \tilde{x}_D < x^*$. Because the social cost function is convex, it is decreasing for $x < x^*$, and therefore, $S_D(\tilde{x}_D) \leq S(\tilde{x})$. Note the stronger result of $S_D(\tilde{x}_D) \leq S(\tilde{x})$ holds for all cases except the trivial case of $\lambda h^* = d = \Delta h = 0$.

5.4.2 When disclosure tax is greater than consumer benefit

The more interesting scenario arises when the disclosure tax is greater than the reduction in consumer harm, because it is no longer clear whether an information disclosure policy reduces social costs. Since consumers often do not pay attention to notices, and do not take actions to decrease their harm, this scenario is also quite likely, as further discussed in Section 6.

First, we know that if $d > \Delta h$, the social cost with disclosure is greater than the social cost without disclosure for any given level of care (i.e., $S_D(x) > S(x)$). However, as shown in Figure 9, the change in social cost evaluated at the firm's optimal level of care, $S_D(\tilde{x}_D) - S(\tilde{x})$, could be greater than, equal to, or less than zero.

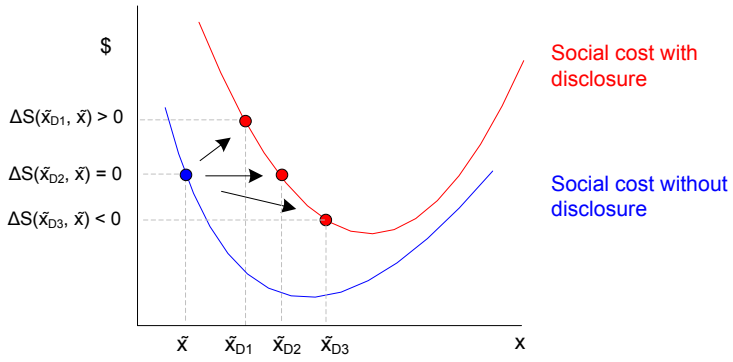

Figure 9: Change in social cost

And so, the critical issue becomes the location on the two curves for which we evaluate ΔS. However, it is clear from the figure that ΔS can be negative, even though $d > \Delta h$. Recall Eq. (12),

$$\Delta S = \Delta \tilde{c} + \tilde{p}_D(d - \Delta h) + \Delta \tilde{p}[i + h] \qquad (12)$$

The first term represents the increased cost due to the disclosure policy and will be positive because *c(x)* is increasing convex and $\tilde{x}_D > \tilde{x}$. The second term is proportional to the difference between the disclosure tax and the reduction in consumer harm which, for $d > \Delta h$, will also be positive. The last term represents the social benefit extracted from the reduced probability of a breach, which is negative because $p(\tilde{x}) > p(\tilde{x}_D)$. And so ΔS will be negative when the last term dominates the first two. This may occur when either (or both) the cost of investigating a breach, *i*, and the total cost of consumer harm without disclosure, *h,* is very large. We also observe that the second term will be smaller when $d - \Delta h$ is very small. Moreover, these effects will be more substantial for smaller levels of care. That is, when both $\tilde{x}_D$ and $\tilde{x}$ are small, the difference in $c(\tilde{x}_D) - c(\tilde{x})$ may well be smaller relative to the change in $p(\tilde{x}_D) - \mathrm{p}(\tilde{x})$ since $c(0) = 0$. This implies that social cost may be reduced by a disclosure regime even if the disclosure tax for firm is substantial compared to the reduction in consumer harm:

Proposition 5*: Even when the disclosure tax is greater than the reduction in consumer harm, an information disclosure policy can still reduce social cost. That is, there* are pairs (d > Δh, λ) such that $S_D(\tilde{x}_D) \leq S(\tilde{x})$.

Proof: By the continuity of $\Delta S(d, \lambda)$ (as shown by Lemma 1) and the fact that $\Delta S(\Delta h, \lambda) < 0$ (Proposition 4), there must exist an interval $I = [\Delta h, \Delta h + \varepsilon)$ such that $\Delta S(d, \lambda) < 0$ if $d \in I$ .The size of the interval, however, depends on the functions *c(x), p(x)* and the specific parameter values *i, h, h** which solves $\Delta S(\Delta h, \varepsilon) = 0$.

Next, we present empirical estimates of the parameters used in this model and provide further discussion.

6. Empirical Validation

Propositions 4 and 5 described the change in social cost when the disclosure tax is less than or equal to, and larger than the reduction in consumer loss (i.e., $d \leq \Delta h$ and $d > \Delta h$). They showed how an information disclosure policy would always be preferred when $d \leq \Delta h$ but that even when $d > \Delta h$, social costs could still be lower. Below, we provide some discussion based on empirical estimates of the parameters in our model, and thus, attempt to determine which case is more likely. As mentioned, we consider *d, h**, and *h* to be measured as 'costs per consumer.' That is, *d* represents the cost of disclosure (tax) to the firm per lost account. Similarly, h_D and *h* represent the average cost per consumer of identity theft (i.e., the actual cost borne by the individual), with and without a disclosure policy.

While robust data is difficult to obtain for *d*, some aggregate estimates are available. For example, Forrester Research claimed a cost of $90 for non-regulated firms and $305 for highly regulated firms (Dignan, 2007). Survey data from the Ponemon Institute over the years 2005-2009 finds the cost per record to be increasing at a decreasing rate: $138 in 2005 and $204 in 2009 (Ponemon, 2010, Figure 1, p12).

Estimates for *h,* range from about \$0-\$300 for median losses and \$422-\$675 for mean losses.[18] Note that these estimates refer to out of pocket expenses and do not include the dollar equivalent of time/effort to address the crime, nor other forms of social loss which may include higher insurance premiums, increased interest rates, civil legal actions, etc. (Baum, 2004). Since these costs represent the loss from all identity theft, we must scale it by the portion of identity theft due to data breaches. Javelin (2006, p3) claims that "businesses as a source of information breach account for 30% of cases" while in a later study they find that only 11% of identity theft is caused by data breaches (Javelin, 2009, Fig. 2). Another study using data from the US Secret Service found that about 26% of identity theft cases were due to data breaches (Gordon et al., 2007). By simply averaging these values, a rough approximation would suggest that data breaches represent about 20% of identity theft giving a potentially more realistic value of $h \approx \$70$.[19]

For the sake of illustration, using these data we obtain estimates of $d \approx \$200$ and $h \approx \$70$. Previously, we showed how information disclosure would always reduce social costs when $d \leq h - h_D$. To our knowledge, the only estimate of h_D comes from Romanosky et al. (2009) who find that data breach disclosure laws reduce identity theft rates by about 2%, on average. If true, this would suggest $h_D \approx \$68.6$ and $\Delta h \approx \$1.4$. It seems clear, however, that even if consumer losses were reduced to zero, the disclosure tax would still be larger than the reduction in consumer harm, $d > h - h_D$.

We can also provide some estimate of the amount of consumer loss borne by the firm, and therefore, the magnitude for λ. Again, robust estimates are difficult to obtain. First, however, if we consider median consumer losses of \$0 as reported by FTC (2007, Table 2) and Javelin (2006, page 2) then this would automatically give $\lambda = 1$. However, note again that these values may likely be underestimates because they do not account for the time and effort involved in addressing the issue (even if the net dollar loss is zero).

If we consider mean (not median) data published by Javelin Research (2006, page 2), out-of-pocket consumer expenses were \$555 (2003), \$675 (2005), \$422 (2006), and total amounts stolen were \$5,249 (2003), \$5,885 (2005) and \$6,383 (2006). If we again consider a 20% portion of loss due to data breaches, this would represent $\lambda = 0.47$ (2003), $\lambda = 0.43$ (2005), and $\lambda = 0.67$ (2006). These results suggest that (at least for the values given above) the firm bears a substantial portion of consumer loss, although obviously, more data is required in order to obtain robust estimates.

As described in Section 3.2, the disclosure tax, *d,* contains many different costs, some of which are incurred directly by the firm (legal and administrative fees, cost of notification, etc.) and others which are imposed by government or industry regulators (e.g., FTC, MasterCard, Visa, etc.). The policy implication is that there is, indeed, some portion of this parameter which is adjustable by a policy maker, and could therefore be manipulated to ensure lower social cost (i.e., that ($d \leq \Delta h$). Empirically distinguishing these costs clearly remains a difficult, yet useful challenge.

In regard to λ, the portion of consumer harm borne by the firm, we recognize that this reflects two main components: the extent to which firms voluntarily (or through social/political norms) provide consumer redress through credit monitoring or identity theft insurance, and the extent to which consumers are successful in bringing individual or class action lawsuits against breached firms. Currently, it appears, however, that such suits are promptly dismissed. For example, negligence claims are largely unsuccessful because plaintiffs are unable to sufficiently

[18] We fully recognize that data is quite sparse, however, the following survey data has been collected: \$0 (FTC, 2006, Table 2, median loss of all forms of identity theft; Javelin, 2006, page 2, median loss), \$500 (FTC, 2003, Table 2, average loss of all forms of identity theft), \$555 (2003), \$675 (2005), \$422 (2006) (Javelin, 2006, page 2, average loss), \$300 (BJS 2005, Table 7, median loss).

[19] We first average the consumer losses, which were approximately \$350 per individual. We then multiply by the scaling factor of 20% which gives \$70.

demonstrate the necessary conditions: causality, actual harm and that the defendant failed to meet a level of due care (Hutchins, 2007; Chandler, 2008; Romanosky and Acquisti, 2009). This suggests that unless firms are somehow induced to directly provide consumer redress through state, federal, or industry sanctions, the current legal tort system remains ill-equipped to fully compensate consumers for harms suffered by data breaches.

From the discussion above, and under the caveat that we are using very limited set of data, it appears that the disclosure tax is substantially greater than the change in consumer harm, yet it would seem that the firm does bear a substantial portion of consumer loss. Given that disclosure is costly, we can presume that (cost-minimizing) firms are increasing their level of care (substantiated by Samuelson, 2007). But it is unclear whether overall social costs have been reduced.

Yet, there may be hope. If it is true that the majority of the disclosure tax is within the control of the firm (i.e., not exogenously imposed) then it is reasonable to assume that the firm will have every incentive to find ways of reducing these costs. In this sense, the firm's incentive is aligned with the social planner. If it is also true that the firm is in the better position to identify and reduce these costs, then this also suggests less demand for government-imposed sanctions (*ex ante* regulation) and more opportunity for a light-handed (paternalistic) policy regime, such as information disclosure.

7. Model Extensions

Below, we extend our basic analysis and briefly consider the effects of a mandatory disclosure policy on consumer action.

7.1 Consumer under-reaction, over-reaction

So far, we have assumed rational consumers who engage in their (privately) optimal level of care, $\tilde{y}_D = y^*$. However, we now relax this assumption to consider consumers who may either under- or over-react to disclosure notifications.

Consumers may (more or less irrationally) under-react when notified of a data breach for a number of reasons. First, as suggested by privacy experts, they may become desensitized to breach notifications causing them to take little or no action (Cate, 2005). Further, certain behavioral decision factors may also cause them to under-react, such as *optimism bias* (consumers perceiving their chances of suffering identity theft to be very low), *rational ignorance* (consumers believing the cost of taking precautions outweighs any benefits they may receive), and *status quo bias* (consumers' own inertia inhibiting them from anticipating the consequences of identity theft and responding) (Romanosky et al., 2008). The effect of these behaviors causes an individual to reduce their incentive to take care, represented as $\tilde{y}_u < y^*$ and $h(y^*) < h(\tilde{y}_u) < h$ as shown in Figure 10. Thus, consumer under-reaction would not qualitatively change any of the results previously described in Section 5.

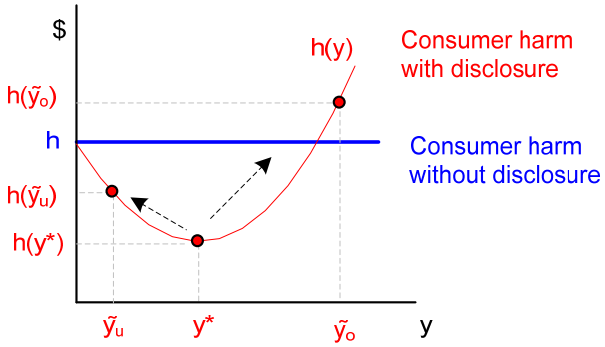

Figure 10: Consumer over- and under-reaction

However, we must also consider that consumers could be driven to over-react, which may increase their expected losses beyond those absent a disclosure regime. For instance, consider the outcome when a consumer receives a breach notification from a bank, closes their account, and opens another account with a competing bank. Quite likely, the consumer has now *increased* their risk of identity theft by disclosing their personal financial information to yet another organization. Similarly, panic, confusion or uncertainty could lead consumers to overestimate the probability of identity theft driving them to purchase identity theft insurance or prevention products which may only increase their expected costs.[20] Indeed, "the hoopla surrounding identity fraud is causing consumers to urgently -- and sometimes blindly -- seek protection" (Wilson, 2007). A recent report on the consumer identity theft prevention market identified more than 20 companies selling various prevention services (Javelin, 2009). Moreover, given recent media attention of the consequences of identity theft, "people are usually willing to pay a premium to protect themselves against the dangers that seem most vivid -- perhaps because they've seen and heard a lot about them" (Regnier, 2005). Such over-reaction outcomes would be represented by $\tilde{y}_o > y^*$ causing harm $h(\tilde{y}_o) > h > h^*$. An important conclusion is that while under-reaction would only limit the benefit of mandatory disclosure, over-reaction could, in fact, increase both the consumer's cost and the overall social cost of disclosure.

The effect of consumer over-reaction is illustrated in Figure 11, where, for simplicity, we plot the change in social cost when the firm internalizes no consumer harm, $\Delta S(d,\lambda = 0)$, for increasing values of consumer harm under mandatory disclosure: $h_D = h$, $h_D = h_1$ and $h < h_D < h_1$ (where $h_1 > h$). All other parameter values are unchanged from Figure 6.

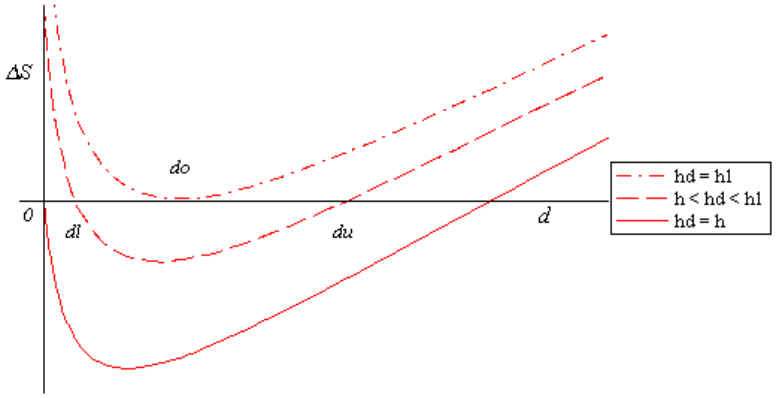

Figure 11: ΔS(d, λ=0) when $h_D > h$

First, observe that the lower curve confirms that the change in social cost equals zero at $d = 0$ when $h_D = h$; i.e., $\Delta S(0,0) = 0$. It is also decreasing when h_D is large: we have $\left.\frac{d\Delta S}{dd}\right|_{d=0,\lambda=0} = \frac{\partial \tilde{x}_D}{\partial d}\frac{\partial \tilde{p}_D}{\partial x}h_D + \tilde{p}_D$. In analogy to the proof of Proposition 3, the first term of this expression is negative and the expression as a whole is negative when h_D is large enough. Therefore, because $\Delta S(0,0) = 0$ and is initially decreasing in d, we see how there are socially beneficial outcomes.

The upper curve, $h_D = h_1$, illustrates how enough consumer over-reaction can eventually prevent a disclosure policy from ever reducing social cost (again, where $\lambda = 0$). Thus, we conclude that for all $h > h_1$, the social cost will always be greater under disclosure, i.e., $\Delta S(d, 0) >$

[20] For example, identity theft insurance is available by a number of leading insurance companies at a cost between $25 and $99 per year (Insure.com, 2009).

0 for all d when $h > h_1$. The value of d at which this minimum is achieved, d_0, is simply the solution to $d\Delta S(d,0)/dd = 0$. We therefore define the value h_1 as the value of h_D that solves $\min \Delta S(d_0,0) = 0$.

Finally, the middle curve, represented by $h < h_D < h_1$, illustrates how the social cost curve intersects the horizontal axis ($d = 0$) now at two locations defined by d_L and d_U. We then state that disclosure will reduce social cost when $d_L < d < d_U$, i.e., $\Delta S(d,0) < 0$ for $d \in [d_L, d_U]$ and $h < h_D < h_1$.

To summarize, under-reaction to a notification can, at worst, only limit consumer harm to a disclosure policy and may still reduce overall social cost. The consequences of over-reaction, however, can be much more severe. At an extreme ($h_D > h_1$), social cost will always be greater under disclosure. Of more interest to a policy maker, however, is a marginal increase in consumer harm ($h < h_D < h_1$) which bounds the amount of disclosure tax at which social cost is reduced. The implication is that even though the reduction in social cost would not be as great as when consumers could take action to reduce their harm ($h_D < h$), social cost could still be lower if a policy maker were able to manipulate the disclosure tax sufficiently, as described in Section 6.

7.2 Moral hazard and second-best social cost

Recall the firm's objective function and social cost function under disclosure,

$$F_D(x) = c(x) + p(x)\left[i + d + \lambda\, h(y)\right] \tag{4}$$

$$S_D(x,y) = c(x) + p(x)\left[i + d + h(y)\right] \tag{6}$$

Eq. (6) is minimized when both the firm and consumer invest in the socially optimal level of care, x^*, y^*, respectively. The firm will invest in x^* when it internalizes the full amount of consumer loss as previously shown.

However, there is an important consequence of this action. Realistically, consumer care (action) is driven by the amount of loss suffered, with level of consumer care decreasing in loss. Lower loss warrants less consumer care and, at an extreme, zero loss warrants zero consumer care. Therefore, if a consumer is completely compensated for their harm, what incentive do they have to take any precaution? The consequence of this moral hazard implies that should the firm invest in x_D^*, the consumer would invest not in $\tilde{y}_D = y^*$ but $\tilde{y}_D = 0$, thereby driving total consumer harm from h^* to h and raising social cost.[21] Therefore, we write

$$S_D(x^*,0) > S_D(x^*,y^*) \tag{21}$$

The policy challenge, therefore, is to induce consumers to take optimal care without having them bear any loss. The tension is this: on one hand, social cost, $S(x,y)$ is minimized when $\lambda = 1$, because this drives $\tilde{x}_D$ to x^*. When the consumer is relieved of all loss, they have no incentive to take care and so we get $S_D(x^*,0) > S_D(x^*,y^*)$. However, in order to achieve the first best level of consumer care, the consumer must suffer all the loss, implying that in order to achieve y^*, λ must be zero. This may suggest that there is an optimal level of λ for which the second-best level of social cost is achieved.

[21] Tort law overcomes this issue by holding victims either partially liable for harm (comparative negligence: each party is liable according to their proportion of fault) or fully liable for harm unless they, themselves take due care (contributory negligence: the injurer pays unless the victim was somehow negligent). Both of these cases induce efficient outcomes by both parties.

7.3 Low cost avoider

A key observation by Ronald Coase (1960) regards the reciprocal nature of social costs (externalities). For example, consider the familiar example of the baker and doctor conducting business in separate offices in the same building (Sturges v. Bridgman, 1879).[22] The doctor complains that the noise from the baker's bread machines drives away his patients and seeks an injunction to prevent the operation of this machine. It is correct that the doctor would not lose patients if the baker stopped baking, but it is also true that the noise would not harm the doctor if he moved to another building. Said differently, the baker imposes a negative externality on the doctor by causing too much noise, however, the doctor would impose an externality on the baker by enjoining him. Therefore, it becomes the joint action of both the doctor and baker operating close by that creates a harm.

Within the context of data breaches: on one hand, consumer information may not have been compromised if the firm had better security controls. But on the other, the consumer's information would not have been stolen if the consumer did not disclose it to the firm.[23] Again, its the joint action between the firm and consumer that creates the loss.

Placed in this context, the social objective, therefore, is to reduce consumer harm at the lowest possible cost. Pigou (1932) might suggest imposing a tax on the firm equal to the consumer harm, thereby inducing them to take the socially optimal amount of care. However, Coase considers that this might be inefficient if the other party (the consumer) can reduce the harm more effectively: "whichever party the blame is assigned to, by government regulators or by the courts, the result will be inefficient if the other party could prevent the problem at a lower cost or if the optimal solution requires precautions by both parties." (Friedman, 2000, p38)

To analyze this, we compare the marginal benefit for each party to reduce consumer harm (i.e., not firm plus consumer harm). Recall the firm's total cost function

$$F_D(x) = c(x) + p(x)\left[i + d + \lambda\, h(y)\right] \tag{22}$$

Which can be rewritten as,

$$F_D(x) = c(x) + p(x)\left[i + d\right] + p(x)\lambda\, h(y) \tag{23}$$

Since we are interested in the cost to the firm of reducing the externality, we represent the firm's marginal benefit of reducing consumer care (only the last term) by,

$$\frac{\partial \hat{F}_D}{\partial x} = p'(x)\lambda\, h(y) \tag{24}$$

Next, recall the consumer's total cost function

$$C_D(x, y) = p(x)\left[1 - \lambda\right]h(y) \tag{25}$$

so the consumer's marginal benefit (of reducing their own care) is given by,

$$\frac{\partial C_D}{\partial y} = p(x)[1 - \lambda]h'(y) \tag{26}$$

[22] Sturges v Bridgman (1879) LR 11 Ch D 852.

[23] Admittedly, it is the case that with databrokers (e.g., Choicepoint) the consumer may not have been involved at all in disclosing their information directly.

Therefore, the firm will be the low cost avoider when

$$\frac{\partial \hat{F}_D}{\partial x} < \frac{\partial C_D}{\partial y} \tag{27}$$

$$p'(x)\lambda h(y) < p(x)[1-\lambda]h'(y) \tag{28}$$

Rearranging gives,

$$\frac{\lambda}{1-\lambda} < \frac{p(x)/p'(x)}{h(y)/h'(y)} \tag{29}$$

Therefore, if the marginal cost (benefit avoided) of the consumer taking action, *y*, to reduce the externality is less than the marginal cost of the firm taking action, *x*, then it is more efficient for the consumer to do so. At some point, however, it will become more costly for the consumer to take action, relative to the firm.

8. Discussion and Limitations

In this article, we constructed an analytical framework to describe the conditions under which a policy of data breach disclosure reduces social costs. Using a methodology commonly employed in the economic analysis of accident (tort) law, we defined cost equations for a firm (injurer; tortfeasor) and a consumer (victim) and illustrated the costs incurred by both parties with and without information disclosure.

Specifically, we showed how mandatory disclosure creates two very important effects. First, it transforms unilateral-care accidents into bilateral-care accidents by enabling both the firm and the consumer to take action to reduce loss. Next, it imposes costs on the firm in two distinct ways. First, the firm will incur direct costs as a result of notification, fines, fees, lost business, etc., what we term the disclosure tax, and are costs a firm would not have incurred but-for public disclosure. Next, we consider that data breach disclosure laws will force the firm to internalize some portion of consumer loss.

We find that both disclosure tax and consumer redress cause the firm to increase its level of care, but only the disclosure tax represents deadweight loss, while redress represents a transfer of costs between the consumer and firm. Therefore, only an increase in redress can reduce the externality caused by the data breach. Further, social cost is always decreasing in consumer redress, but if this is small enough, some disclosure tax is necessary to reduce social cost. Therefore, if the firm bears only a small portion of consumer harm, the social planner may be justified in applying (or threatening to apply) additional fines or fees on the firm in order to minimize social cost.

At central issue is the comparison between the disclosure tax and the benefit achieved from reduced consumer harm (identity theft). We show that mandatory disclosure is always preferred when the disclosure tax is less than, or equal to, the benefit from lower consumer harm. However, even when the disclosure tax is greater than the benefits, social costs may still be lower.

While we believe we have addressed the key issues in this work, we describe a number of limitations and alternative approaches below.

Super-Consumer and Attacker Models

As mentioned above, information disclosure enables consumers to punish firms for the firm's bad behavior. For instance consider a retail bank that suffers a data breach and a consumer

who punishes the firm by changing banks. In such a case, a researcher may consider modeling both the firm that lost a customer and the firm that gained a customer. However, one firm's loss is another's gain (zero-sum) and therefore the net social impact would only be the consumer's switching cost which would not materially affecting our conclusions. Therefore, we have restricted the analysis simply to the breached firm and affected consumer.

Moreover, one may consider many consumers affected by a breach, instead of just one. If we define the consumer cost function 'per-consumer' then one would only need to multiply the cost function by N, the number of consumers affected by a breach. However, again, this would not materially change our results. Alternatively, we could just define the consumer cost equation as "total consumer loss." Consider another example: when thieves use stolen credit card information to purchase goods, the retail merchant where the card is used may bear the cost of the fraudulent purchase (considered a "card-not-present" transaction). In general, considering the other parties affected by a single firm's breach simply leads to the notion of a *super-consumer* that incurs some harm as a result of a data breach and our model remains mathematically unchanged, and only differs in interpretation. Again, we simplify the analysis by considering just one consumer.

In cases where information is stolen, not lost, one may consider modeling the attacker's costs and motivation as a function of increased firm care. There are two reasons why we have not pursued this here (though certainly a relevant question). First, hacking tools are quite often scripted and fully automated. Therefore, once functional exploit code exists, the marginal cost of launching a new attack is close to zero. Second, if most security incidents are initiated from countries other than from where the firm and consumer reside, then one might question the justification for including attacker costs or benefits in the social cost model.

Alternative Methodologies

In lieu of our methodological approach, there are other possible techniques one may use to address these and other issues. Our research question might be resolved using a Stackelberg game of strategic substitutes (Miceli, 1997, p59): the more the firm invests in care, the less the consumer need invest (here, quantity of some good is replaced by level of care). This also captures the sequential nature of data breaches and resulting consumer harm and the possibility that suboptimal investment in care by one party causes overinvestment by another.

Alternatively one might consider a Hotelling model that describes two firms competing for one consumer and each firm's investment in security controls (and their willingness to disclose breach information) reflects the horizontal differentiation exploited by such models. A researcher may also consider a competitive firm model employing information asymmetry and that a consumer's decision to purchase from one firm depends on their subjective probability of that firm suffering a breach contemporaneously conditional on the firm not having previously suffered a breach. Further, one might introduce variation by modeling the strategic decisions of consumers who vary in their privacy preferences: fundamentalists, pragmatists and unconcerned.

Level of Care vs Level of Activity

The most common decision variable in economic analyses of law is the level of care by either one or both parties. Extensions to this work include *level of activity* as a means of reducing total cost. That is, one may consider that the total cost of accidents is not only a function of one's level of care, but in the amount of activity in which they engage. For instance, total social costs from car accidents may be a function of the speed at which one drives (level of care) and the number of miles they drive (level of activity). A number of interesting research questions follow: i) is it better to reduce the consumer's or firm's levels of activity, and which approach produces the largest marginal reduction in cost or increase in benefit?; ii) is it better to reduce level of activity or increase care?; iii) is it better to increase firm or consumer care?

Positive Externalities of Increased Security Investment

Improvements in IT security imply that the organization is better able to prevent the compromise of corporate (employee, customer, trade secrets) data, is less vulnerable to unauthorized modification of IT data, and is more resilient to system outages or degraded performance. Together these three approaches represent the familiar IT security control landscape of confidentiality, integrity and availability (C, I, A). As a consequence of this improvement, an organization may realize benefits not only from fewer data breaches, but also from the ability to avoid other kinds of security incidents, such as a computer virus manipulating or destroying employee data, or extended outages of production business services. Therefore, a positive externality of this sort may reduce the firm's cost of a disclosure law as a function of firm's care.

9. References

Arora, A., Telang. R, Xu, H. (2007). *Optimal Policy for Software Vulnerability Disclosure*. Management Science, 54(4), 642-656.

Baum. K. (2004). *Identity Theft, 2004*. Bureau of Justice Statistics.

Baum. K. (2005). *Identity Theft, 2005*. Bureau of Justice Statistics.

Beales, H., Craswell, R., Salop, S. (1981). *The Efficient Regulation of Consumer Information*. Journal of Law and Economics, 24(3).

Beebe, M. (2009). *Missouri Becomes the 45th State to Enact Data Breach Notification Legislation.* Perkin Coie Digestible Law. Available at http://www.digestiblelaw.com/datasecurity/blogQ.aspx?entry=6064&id=34.

Best, J. (2009). *HSBC companies slapped with US$5M fines over data breaches*. ZDNet Asia. Available at http://www.zdnetasia.com/news/business/0,39044229,62056295,00.htm.

Brodkin, J. (2007). *ChoicePoint settles with 43 states over data breach*. Network World. Available at http://www.networkworld.com/news/2007/053107-choicepoint-settles-data-breach.html.

Brown, J. (1973), *Toward an Economic Theory of Liability*, J. of Legal Studies, 2(2), 323-349.

Buckley, W. and Okrent, C., (2003). *Torts & Personal Injury Law (3rd ed)*. Delmar Cengage Learning, 46.

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

Cate, F. (2005, February 27). *Another notice isn't answer*. USA Today. Available at http://www.usatoday.com/news/opinion/2005-02-27-consumer-protection-oppose_x .htm. Access date October 15, 2009.

Cavusoglu, H., Cavusoglu, H. and Zhang, H. (2008). *Security Patch Management: Share the Burden or Share the Damage?* Management Science, 54(4), 657-70.

Cavusoglu, H., Mishra, B. and Raghunathan, S. (2004). *The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers*. International Journal of Electronic Commerce, 9(1).

Chandler, J. (2008). *Negligence Liability for Breaches of Data Security*, Banking and Financial Law Review, 23, 223-47.

Coase, R. (1960). *The Problem of Social Cost*. Journal of Law and Economics, 3, 1-44.

Dignan, L. (2007, May 8th). *What that data breach will really cost you. ZDNet*. Referencing Figure 1 reproduced from Forrester Research. Available at http://blogs.zdnet.com/BTL/?p=5007.

Federal Trade Commission. (2003). *2006 Identity Theft Survey Report*.

Federal Trade Commission. (2007). *2006 Identity Theft Survey Report*.

Friedman, D. (2000). *Law's Order: What Economics Has To Do With It And Why It Matters.* Princeton University Press.

Fishman, M. and Hagerty, K. (2003). *Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers*. Journal of Law, Economics and Organization, Oxford University Press, 19(1), 45-63.

Fung, A., Graham, M., & Weil, D. (2007). *Full Disclosure: The Perils of and Promise of Transparency*. Cambridge University Press.

GAO. (2007). *Data Breaches Are Frequent, but Evidence of Resulting Identity Theft Is Limited; However, the Full Extent Is Unknown*. Government Accountability Office. GAO-07-737.

Gandal, N., Choi, J. and Fershtman, C. (forthcoming 2009). *Network Security: Vulnerabilities and Disclosure Policy*. Journal of Industrial Economics.

Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., and Konstan., J. (2005). *Stopping Spyware at the Gate: A User Study of Privacy, Notice and Spyware*. Symposium on Usable Privacy and Security (SOUPS), Carnegie Mellon University.

Goodin, D. (2008). *IT Contractor Caught Stealing Shell Oil Employee Info*. The RegisterUK. Available at http://www.theregister.co.uk/2008/10/07/shell_oil_database_breach/.

Goodin, D. (2009). *Data-sniffing attack costs Heartland $12.6m*. The Register UK. Available at http://www.theregister.co.uk/2009/05/07/heartland_breach_costs/.

Gordon, G., Rebovich, J., Choo, K., Gordon, J. (2007). *Identity Fraud Trends and Patterns: Building a Data-Based Foundation for Proactive Enforcement*. Center for Identity Management and Information Protection, Utica College.

Grossklags, J., Christin, N., and Chuang, J. (2008). *Secure or Insure? A Game-Theoretic Analysis of Information Security Games*. In Proceedings of the 17th International World Wide Web Conference, Beijing, China, 209-218.

Hogan, M. (2008). *Arrests Made in ID Theft Case*, Sealy News. Available at http://www.sealynews.com/articles/2008/09/23/news/news04.prt.

Hutchins, J. (2008, August 8). *A New Frontier in Privacy Litigation: The Advent of Private Lawsuits Over Data Security Breaches*. Remarks at the ABA Annual Meeting.

Insure.com. (2009, July). *Who's who in the identity theft protection market*. Available at http://www.insure.com/articles/idtheft/id-theft-products.html.

Javelin Strategy and Research. (2006). *Identity Fraud Survey Report: Consumer Version*.

Javelin Strategy and Research. (2009). *Identity Fraud Survey Report: Consumer Version*.

Javelin Strategy and Research. (2009). *Consumer Identity Protection Services Scorecard: Competition Intensifies as Vendors Aggressively Expand Offerings*.

Jin G. and Leslie, P. (2003). *The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards*. Quarterly Journal of Economics, 118(2), 409-51.

Kaplan, D. (2009). *TJX settles for $525K with four banks over breach*. SC Magazine. Available at http://www.scmagazineus.com/TJX-settles-for-525K-with-four-banks-over-breach/article/148095/.

Kannan, K., Rees, J., and Sridhar, S. (2007). *Market Reactions to Information Security Breach Announcements*. International Journal of Electronic Commerce, 12 (1).

Kerber, R. (2007). *Cost of data breach at TJX soars to $256m*. The Boston Globe. Available at http://www.boston.com/business/globe/articles/2007/08/15/cost_of_data_breach_at_tjx_soars_to_256m/.

Kravets, D. (2008). *Judge Weighing Ameritrade Hack Lawsuit Settlement*. Wired Magazine. Available at http://www.wired.com/threatlevel/2008/06/judge-weighing/.

Krebs, B. (2008). *Thieves Stole Identities to Tap Home Equity*. Washington Post.

Kolstad, C., Ulen, T. and Johnson, G. (1990). *Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements?* American Economic Review, American Economic Association, vol. 80(4), 888-901.

Landes, W. and Posner, R. (1987). *The Economic Structure of Tort Law*. Harvard University Press.

Lenard, T. and Rubin, P. (2005). *Slow Down on Data Security Legislation*. Progress Snapshot 1.9. The Progress & Freedom Foundation.

Lemos, R. (2009). *TJX estimates breach costs at $118 million*. Security Focus. Available at http://www.securityfocus.com/brief/568, accessed 02/28/09.

Maurushat, A. (2009). *Data Breach Notification Law Across the World from California to Australia*, University of New South Wales Faculty of Law Research Series, http://law.bepress.com/cgi/viewcontent.cgi?article=1153&context=unswwps.

McMillan, R. (2008). *United Healthcare Data Breach Leads to ID Theft*. Network World.

McMillian, R. (2009, February 23). *Starbucks sued after laptop data breach*. Network World. Available at http://www.networkworld.com/news/2009/022309-starbucks-sued-after-laptop-data.html. Accessed November 20, 2009.

McGlasson, L. (2009, February 27). *Heartland Data Breach: Class Action Suit Filed on Behalf of Banking Institutions*. BankInfoSecurity.com. Available at http://www.bankinfosecurity.com/articles.php?art_id=1239. Accessed November 20, 2009.

Miceli, T. (1997). *Economics of the Law: Torts, Contracts, Property, Litigation*. Oxford University Press.

Mulligan, D. (2007). *Information Disclosure as a light-weight regulatory mechanism*. Presentation at DIMACS, Workshop on Information Security Economics, Rutgers University.

Pigou, A. (1932). *The Economics of Welfare*. Library of Economics and Liberty, 22.

Polinsky, M. and Shavell, S. (2006). *Mandatory versus Voluntary Disclosure of Product Risks*. Stanford Law and Economics Olin Working Paper No. 327.

Ponemon Institute. (2010). 2009 *Annual Study: Cost of a Data Breach*. The Ponemon Institute.

Pulliam, D. (2007). *VA sets aside $20 million to handle latest data breach*. Government Executive. Available at http://www.govexec.com/story_page.cfm?articleid=37191&dcn=todaysnews.

Regnier, P. (August 22, 2005). *The ID theft protection racket: Are you terrified about identity theft yet? If not, consider this: It could get you killed*. MONEY Magazine. Available at http://money.cnn.com/2005/08/22/pf/idtheft_0509/

Romanosky, S., Telang, R., and Acquisti, S. (2008). *Do Data Breach Disclosure Laws Reduce Identity Theft?* Available at SSRN: http://ssrn.com/abstract=1268926.

Romanosky, S. and Acquisti, A. (2009). *Privacy Costs and Personal Data Protection: Economic and Legal Perspectives of Ex Ante Regulation, Ex Post Liability and Information Disclosure*, 24 Berkeley Technology Law Journal.

Samuelson Law, Technology, & Public Policy Clinic. (2007). *Security Breach Notification Laws: Views from Chief Security Officers*. University of California-Berkeley School of Law.

Schwarzenegger, A. (2007). Letter to the members of the California State Assembly, available at http://gov.ca.gov/pdf/press/2007bills/AB%20779%20Veto%20Message.pdf.

Simon, J. (2009). *States' weapon of choice against ID theft: Transparency*. Creditcards.com. Available at http://www.creditcards.com/credit-card-news/data-security-breach-notification-laws-1282.php.

Shavell, S. (1984). *A Model of the Optimal Use of Liability and Safety Regulation*, RAND Journal of Economics, The RAND Corporation, vol. 15(2), 271-280.

Shavell, S. (2004). *Foundations of Economic Analysis of Law*. Harvard University Press.

Telang, R. and Wattal, S. (2007). *An Empirical Analysis of the Impact of Software Vulnerability Announcements on Firm Stock Price.* IEEE Transactions on Software Engineering paper, 33 (8), 544-57.

Wang, T., Rees, J., and Kannan, K. (2009). *The Impact of Information Security Disclosures on Market Reactions to Security Breaches*. (under review).

Wilson, T. (Dec 19, 2007). *Amid Confusion, Market for ID Theft Services Grows.* DarkReading.com. Available at http://www.darkreading.com/security/vulnerabilities/showArticle.jhtml?articleID=208803786.

Zetter, K. (2009). *Judge Revives Question of Retail Liability in Hannaford Breach Case.* Wired Magazine. Available at http://www.wired.com/threatlevel/2009/10/hannaford/.

Zick, C. (2009). *California Hospital Fined $187,500 For Octuplet Mom Breach*. SecurityPrivacyLaw Blog. Available at http://www.securityprivacyandthelaw.com/2009/07/articles/data-breach-1/california-hospital-fined-187500-for-octuplet-mom-breach/.

10. Appendix

Lemma 1: Continuity of ΔS(d, λ)

ΔS is a continuous function $\forall (d,\lambda) \in [0,\infty)$ *and* $[0,1]$. This follows from the existence and uniqueness of $\tilde{x}_D$ given *d* and λ, which in turn is a consequence of the assumptions on the functions *c(x)* and *p(x)*.

Lemma 2: Social costs bounded between ΔS(d, λ = 0) and ΔS(d, λ = 1)

For all *d*, $\Delta S(d,1) < \Delta S(d,\lambda) < \Delta S(d,0)$. We can prove this by showing that ΔS is decreasing in λ for all *d*. Consider the following:

$$\frac{d\Delta S}{d\lambda} = \frac{d\Delta\tilde{c}}{d\lambda} + \frac{d\Delta\tilde{p}}{d\lambda}(i+h) - \frac{d\tilde{p}_D}{d\lambda}(d-\Delta h) \tag{30}$$

$$\frac{d\Delta S}{d\lambda} = \left(\frac{\partial\tilde{c}_D}{\partial x} + \frac{\partial\tilde{p}_D}{\partial x}(i+d+h^*)\right)\frac{\partial\tilde{x}_D}{\partial\lambda} \tag{31}$$

$$\frac{d\Delta S}{d\lambda} = \left(-\frac{\partial\tilde{p}_D}{\partial x}(i+d+\lambda h^*) + \frac{\partial\tilde{p}_D}{\partial x}(i+d+h^*)\right)\frac{\partial\tilde{x}_D}{\partial\lambda} \tag{32}$$

$$\frac{d\Delta S}{d\lambda} = \frac{\partial\tilde{p}_D}{\partial x}\frac{\partial\tilde{x}_D}{\partial\lambda}(1-\lambda)h^* \tag{33}$$

Note that Eq. (33) follows from the identity,

$$\frac{\partial \tilde{c}_D}{\partial x} = -\frac{\partial \tilde{p}_D}{\partial x}(i + d + \lambda h^*), \tag{34}$$

the condition that must be satisfied in order for $\tilde{x}_D$ to be the argument which minimizes the firm's costs under disclosure. This identity also allows us to determine $\partial \tilde{x}_D / \partial \lambda$,

$$\frac{\partial}{\partial \lambda}\frac{\partial \tilde{c}_D}{\partial x} = -\frac{\partial}{\partial \lambda}\frac{\partial \tilde{p}_D}{\partial x}(i + d + \lambda h^*) \tag{35}$$

$$\frac{\partial^2 \tilde{c}_D}{\partial x^2}\frac{\partial \tilde{x}_D}{\partial \lambda} = -\frac{\partial^2 \tilde{p}_D}{\partial x^2}\frac{\partial \tilde{x}_D}{\partial \lambda}(i + d + \lambda h^*) - h^*\frac{\partial \tilde{p}_D}{\partial x} \tag{36}$$

$$\frac{\partial \tilde{x}_D}{\partial \lambda} = -\frac{h^*\dfrac{\partial \tilde{p}_D}{\partial x}}{\dfrac{\partial^2 \tilde{c}_D}{\partial x^2} + \dfrac{\partial^2 \tilde{p}_D}{\partial x^2}(i + d + \lambda h^*)} \tag{37}$$

$$\frac{\partial \tilde{x}_D}{\partial \lambda} > 0 \tag{38}$$

Where the last line follows from the monotonicity of p*(x)* and the strict convexity of p*(x)* and *c(x)*. Eq. (30) then implies that

$$\frac{d\Delta S}{d\lambda} < 0 \tag{39}$$

And therefore ΔS is decreasing in λ for all *d*.

Lemma 3: ΔS tends to infinity as *d* increases

To show that $\lim_{d\to\infty} \Delta S = \infty$, we show that $\lim_{d\to\infty} c(\tilde{x}_D) = \infty$ (i.e., Δc and thus ΔS both diverge to infinity). From the equation for $\tilde{x}_D$, we see that $\lim_{d\to\infty} -c'(x)/p'(x) = \lim_{d\to\infty} i + d + \lambda h^* = \infty$.

Therefore either $c'(\tilde{x}_D) \to \infty$ or $p'(\tilde{x}_D) \to 0$ as $d \to \infty$. If $\lim_{d\to\infty} p'(\tilde{x}_D) = 0$, then $\tilde{x}_D$ must tend to infinity (otherwise the convexity of *p(x)* would be violated). In this case, the convexity of *c(x)* implies that $\lim_{d\to\infty} c(\tilde{x}_D) = \infty$.

If $\lim_{d\to\infty} c'(\tilde{x}_D) = \infty$, then either $\tilde{x}_D \to \infty$ (in which case $c(\tilde{x}_D) \to \infty$) or $\tilde{x}_D$ tends to some finite value, $\hat{x}$. Then by the continuity of *c(x)* and its derivatives, $\lim_{\tilde{x}_D\to\hat{x}} c'(\tilde{x}_D) = \infty$, and thus $\lim_{d\to\infty} c(\tilde{x}_D) = \lim_{\tilde{x}_D\to\hat{x}} c(\tilde{x}_D) = \infty$, by the convexity of *c(x)*.

Lemma 4: ΔS(d,0) and ΔS(d,1) converge as *d* approaches infinity if c(x) is strongly convex

Given that λ does not appear explicitly in the expression for Δ*S*, it is sufficient to show that $\tilde{x}_D(d,\lambda=0) \to \tilde{x}_D(d,\lambda=1)$ as $d \to \infty$. We can write the first order Taylor expansion for $\tilde{x}_D(d,\lambda)$ in λ as

$$\tilde{x}_D(d,1) = \tilde{x}_D(d,0) + \frac{\partial \tilde{x}_D}{\partial \lambda}(\xi - 0) \tag{40}$$

$$\tilde{x}_D(d,0) - \tilde{x}_D(d,1) = -\xi \frac{\partial \tilde{x}_D}{\partial \lambda} \tag{41}$$

$$\left| \tilde{x}_D(d,0) - \tilde{x}_D(d,1) \right| \leq \left| \frac{h^* \frac{\partial \tilde{p}_D}{\partial x}}{\frac{\partial^2 \tilde{c}_D}{\partial x^2} + \frac{\partial^2 \tilde{p}_D}{\partial x^2}(i + d + \lambda h^*)} \right| \tag{42}$$

The right hand side of Eq. (42) tends to zero as $d \to \infty$ because either $\tilde{x}_D \to \infty$ or $\tilde{x}_D \to \hat{x}$. If $\tilde{x}_D \to \infty$, the numerator approaches zero by the properties of *p(x)*, and the denominator is bounded away from zero if *c(x)* is strongly convex. Note that this is a slightly stronger condition than the strict convexity which has been assumed so far. If $\tilde{x}_D \to \hat{x}$, then $\frac{\partial^2 \tilde{p}_D}{\partial x^2} \to \frac{\partial^2 \hat{p}}{\partial x^2} > \varepsilon > 0$ for some ε, and the right hand side tends to zero as 1/*d*.

Lemma 5: There exists a value of *d* such that ΔS(d, λ) = 0 for $h_D < h$

Given that the continuity of ΔS (Lemma 1) and Lemma 3, there must be a point such that Δ*S* = 0. Moreover, because $d\Delta S / d\lambda < 0$ (Lemma 2), then $\Delta S(0,\lambda) < 0 \ \forall \ \lambda$. Let us define d_u as the value of *d* such that ΔS = 0 when $h_D < h$. That is, d_u is the implicit solution to Eq. (10),

$$\Delta S = \Delta \tilde{c} + \tilde{p}_D(d - \Delta h) + \Delta \tilde{p}[i + h] = 0 \tag{12}$$

$$d_u = \frac{\Delta h - \Delta \tilde{c} - \Delta \tilde{p}[i + h]}{\tilde{p}_D} \tag{43}$$

The solution is implicit because $\Delta \tilde{c}$, $\Delta \tilde{p}$ and $\tilde{p}_D$ are also functions of *d*.