Data Security in the Digital Age: Reputation and Strategic Interactions in Security Investments

Ying Lei Toh[*]
Toulouse School of Economics

(Preliminary draft. Please do not cite or circulate.)

May 10, 2016

Abstract

This paper develops a model of data security investment in which concerns about reputation - consumers' belief about a firm's security level - provide the underlying incentive for the firm to invest. I consider two set-ups. In baseline model, a website makes an unobserved one-time security investment that consumers learn about over time. Underinvestment in security arises as the website does not internalise consumers' losses from data breaches; imperfect breach detection further limits the consumers' ability to punish the firm. In the extended model, I introduce the consumers' bank. Under this set-up, the overall level of security is jointly determined by the website's and the bank's investment levels. I examine the implications of regulatory policies - specifically, mandatory breach notification and a minimum security standard - targeted at raising the level of security. I show that these well-intentioned policies may not always result in a higher overall level of security, demonstrating the importance of accounting for the agents' strategic behaviours in regulatory interventions.

1 Introduction

Rapid advancements in technology have enabled firms to collect, analyse and store vast amounts of consumer information in an unprecedented fashion. While this has created

[*] First and foremost, I would like to thank my thesis advisor, Bruno Jullien for his patient guidance and support. I would also like to extend my gratitude towards Paul Seabright, Jacques Crémer, Michael Riordan, Alexandre de Cornière, Rossella Argenziano, Ananya Sen, Holger Herbst, Miaomiao Dong, Simon Fuchs, Alexia Gonzalez, the members of the applied theory group at TSE, the participants at the ENTER seminar at the University of Mannheim and the 2016 Bi-annual Postal Economics Conference and the anonymous reviewers of WEIS 2016 for their helpful feedback and comments.

many benefits for consumers (such as the personalisation of products and services), it has also sparked concerns about data security. Consumer data collected by firms is susceptible to data breaches, which may arise from either within (e.g. from the loss of data containing devices or acts of malicious insiders) or without (e.g. via hacking) the organisation. Target, Home Depot, Ebay, JPMorgan Chase, and Ashley Madison are a few recent examples of the many firms that have been affected by massive data breaches. In 2015, the number of reported data breach incidents reached an all-time high - a total of 3930 incidents, resulting in over 736 million compromised records (Risk Based Security (2015)). As more businesses go digital and as cybercriminals become more sophisticated, the number of data breaches can be expected to continue to grow.

Data breaches create negative repercussions that often extend beyond the affected firms. Data breaches involving consumer data, for example, can result in identity thefts, phishing scams and payment card frauds; the associated costs of which are typically borne by consumers and other agents (e.g. banks or payment card issuers) in the economy. In addition to the immediate costs that data breaches impose, their prevalence may erode consumers' trust and, ultimately, impede the development of the digital economy. It is thus unsurprising that data protection – from the reform of the Data Protection Directive in the EU to the policy discussions surrounding the proposed data breach notification bills in the U.S. – is high on the agenda of policy makers around the world.

One reason for the pervasiveness of data breaches may be that firms are underinvesting in security. For any given level of security threat, a firm decides on its optimal level of investment by weighing the benefit from reducing its vulnerability to attacks against the cost of investment (Gordon and Loeb (2002)). The benefit of investment may be decreased in the presence of informational asymmetries and externalities, resulting in low levels of investment (Anderson and Moore (2006)). Externalities in security investments may exist between firms. With a growing level of interconnectedness between firms, an individual firm's security level becomes increasingly dependent on the effort/investment levels of the other firms in its network. Some studies have attempted to capture the interdependencies in security investments in economic frameworks, either in public good games (Varian (2004), Grossklags et al. (2008)) or in independent security games (Kunreuther and Heal (2003)). Consider first security investment as a public good. This is typically illustrated with the analogy of a city that is protected by a wall, whose strength depends on the effort exerted by the individual builders (Varian (2004)). Varian (2004) considers three different types of public good games - total effort, weakest link and best shot - and Grossklags et al. (2008) introduce a fourth - weakest-target. Because of the public good nature of security, free-riding occurs in equilibrium (in most of the aforementioned classes of games), implying an underinvestment in security. Interdependencies can also arise due to the potential contagion of security attacks. Kunreuther and Heal (2003) show that, in interdependent security games, the contagion effect reduces an individual's benefit from investing in security. This results in a sub-optimal level of investment in equilibrium.

Informational asymmetries and externalities present between a firm and its consumers are also a potential cause of the underinvestment in security. First, unless made liable, firms do not internalise the costs that data breaches impose on consumers (and the society at large) when making their investment decision. Thus, the firm underinvests relative to the social optimum. Second, security investments and data breaches are typically unobserved by the consumers, which implies that firms are not sufficiently rewarded (or punished) for their investment (or the lack thereof) in security. This may lead to a moral hazard problem. In this regard, investment in data security bears much resemblance to investment in unobserved product quality. In economic models of unobserved endogenous quality choice, reputation – defined as the consumers' belief about the firm's quality – is usually considered to be the main driver of a firm's investment (Shapiro (1983), Board and Meyer-ter Vehn (2013)). Consumers in these models receive signals about quality, which allows them to update their beliefs, leading to reputation adjustments over time. In the context of data security, one can think of data breaches as a form of quality (or security) signal. Empirical studies (e.g. Campbell et al. (2003), Cavusoglu et al. (2004), Acquisti and Grossklags (2005)) have demonstrated that breach announcements can have a negative impact on the firm's stock price or market valuation. To the extent that the investors' reaction reflects the consumers', reputation effect may be an important motivating factor for security investment. One may expect the effect of reputation on security investment to be small, however, since consumers are unlikely to learn about data breaches unless they are disclosed by the firm. Consequently, firms would have little incentive to invest in security.

The preceding discussion illustrates how externalities and information asymmetries – between interconnected firms and also between a firm and its consumers – may help to explain the firms' lack of incentives to invest in security. Existing models of data security (as presented earlier) have attempted to capture the interdependencies between firms; however, they largely focus on the defence of a network and do not explicitly model the consumer side of the market. The loss from data breaches for firms are typically taken to be exogenously given. This paper aims to fill this gap in the literature. I begin by developing a model of unobserved security investment, in which reputation concerns motivate the firm to invest in security. To my knowledge, this is the first paper to model data security as an unobserved quality dimension of the firm. Similar to what is commonly done in the unobserved quality literature, I define reputation as the consumers' belief about the website's level of security and consider the detection of a data breaches as bad news signals. In modelling the data security problem, I adopt the probability-based approach commonly used in the information security literature. More specifically, I model security threat as the probability of a data breach; this probability may be reduced via the firm's investment in security. I further extend the model to allow for interdependencies in security investments.

The baseline model in this paper focuses on the interaction between the firm and its consumers. The set-up consists of an on-line retailer (that I will refer to as the website) and a unit mass of consumers who live for two periods. The consumers derive utility - which is heterogeneous across the population - from the website's product, but suffer a loss when the payment card information they provide to the website is breached. The probability of a breach is determined by a one-time security investment made by the website at the beginning of the game. Neither the amount of investment nor its outcome is observed by the consumers. In each period, the consumers have to decide whether to use the website given their valuation for its product and their belief about its level of security (i.e. the website's reputation for security). If they use the website and a breach occurs, they learn about it with some probability. The detection of a breach leads to an unfavourable updating of beliefs (i.e. a poorer reputation for security); this generates customer turnover and revenue loss for the website. It is this reputation effect of a data breach that provides the website with incentive to invest in data security.

Underinvestment in security arises because the website does not internalise the cost of data breaches to consumers. Although reputation concerns do provide the website with some incentive to invest, the effect may be weak as data breaches often escape undetected by consumers. One policy that can help to strengthen the reputation effect is mandatory breach notification. I show that when the consumers' loss from data breaches are treated as exogenous, mandatory breach notification does indeed raise the website's level of security investment. I then introduce the possibility of consumer self-insurance. The interaction between protection and insurance in security investment has also been analysed in Grossklags et al. (2008). They study firms' choice of self-protection and self-insurance strategies in a public good game; both strategies are decided at the beginning of the game. The nature of the interaction I consider differs from that of Grossklags et al. (2008). In my model, security is a private good and the firm is not able to self-insure. Instead, I study how the incentive of the firm to self-protect changes when consumers are able to self-insure. Further, the level of self-insurance is taken to be exogenously given. Returning to the earlier discussion, I show that breach notification may not necessarily make the consumers better off when they can self-insure. Breach notification enables consumers to take actions to mitigate a fraction of the loss. This lowers the consumers' expected loss from using an insecure website, which in turn reduces the consumer turnover when a breach occurs. In other words, the effect of self-insurance countervails the reputation effect of breach notification. I find that when the ability of consumer to self-insure is high, mandatory breach notification leads to a decrease in the level of security investment made by the website. This serves as a first illustration of the importance of accounting for strategic interactions between stakeholders in policy analysis. My finding lies in contrast with that of Romanosky et al. (2010), who look at how mandatory breach notification affects the levels of firms' security investment and consumers' self-insurance (which they refer to as consumer care). Because they assume that breach notification imposes an exogenous disclosure cost that is not affected by consumer self-insurance, the firm always invests more under breach notification.

In the extended model, I bring in the consumers' bank by interpreting consumer self-insurance to be the fraud liability assumed by the bank. Put differently, when a consumer reports a fraudulent transaction made on his card to his bank, the bank reimburses him a fraction of the loss. I consider the situation where the probability of a fraud depends on both the website's and the bank's security investments, introducing interdependencies into the framework. Unlike in existing works on information security where the players contribute to a single line of defence, in this paper, the website's and the bank's protection form two lines of defence. The website's security determines the probability that information will be breached and the bank's security influences the likelihood that fraud can be committed with the stolen information. I assume that the consumers observe the bank's security level. Further, they can only learn about a data breach when it results in fraud; that is, when both lines of defence are breached. I show that the bank's security investment is always a strategic substitute for the website's (i.e. the website's investment is decreasing with the bank's), whereas the website's investment may either be a strategic complement or substitute to the bank's, depending on the type of bank security measures considered. The bank's security level may also affect consumer behaviour. A high level of security provided by the bank may encourage consumers to use an insecure website, creating a consumer moral hazard problem. Should this arise, the website would have no incentive to invest. Multiple equilibrium outcomes are possible in this framework. In particular, a unique and stable Nash equilibrium, in which both the website and the bank invest in security, exists when the bank's marginal cost of investment is sufficiently high. Taking this equilibrium as a starting point, I examine how a minimum bank security standard and mandatory breach notification may affect the website's and bank's incentives to invest. I demonstrate in both cases that the policies may not necessarily lead to a higher overall level of security. This again highlights the importance of considering strategic behaviours policy design and evaluation.

This paper is structured as follows. Sections 2 presents the set-up of the baseline model. Section 3 and 4 provide the equilibrium analysis and the policy analysis of a mandatory breach notification law respectively. Section 5 extends the baseline model to include the consumer's bank and examines the impact of a minimum bank security standard and mandatory breach notification. Section 6 concludes. All proofs will be deferred to the mathematical appendix.

2 Model Set-up

Consider a two-period model with a website and a unit mass of consumers. The website sells a product for which the consumers value heterogeneously; the consumers' valuation, v, is uniformly distributed between 0 and 1. When the consumers purchase the product, a revenue of r is generated for the website. As part of the transaction, information about the consumers is also collected and stored by the website. This information may be compromised by hackers, generating a loss of L to the consumers.

Let ρ denote the probability of a breach. I consider two possible states of security - "good" and "bad" - with corresponding probabilities of breach ρ_G and ρ_B. The probability that a breach occurs in the "good" state is lower than that of the "bad" state; i.e., $0 \leq \rho_G < \rho_B \leq 1$. For simplicity, let us set $\rho_G = 0$. In this paper, I will also refer to the a website in a "good" state of security as a secure website and that in a "bad" state as an insecure website. The state of security depends on the website's investment in data security. Prior to interacting with the consumers, the website makes a one-time security investment of $c(q^f)$, where q^f is the probability that the "good" state is achieved. The state of security remains the same across the two periods. One can think of the "bad" state of security as corresponding to the situation where the website's security investment is rendered ineffective in the future periods. For example, cybercriminals may manage to develop a tool to circumvent the security measure(s) that the firm has invested in. In this case, $1 - q^f$ gives the probability that such a tool is developed and ρ_B denotes the probability that the cybercriminals succeed in stealing consumer information with the tool. Figure 1 illustrates the security outcomes for a given level of investment c.

The consumers observe neither the amount nor the outcome of investment; however, they form rational expectations over the expected probability of a data breach. I define the website's reputation for security to be the consumers' belief that the "good" state is achieved. Let q^c denote the consumers' initial belief that the "good" state of security is attained (or the initial reputation of the website); i.e. $Pr(\rho = \rho_G) = q^c$. Although the true state of security is not observed by the consumers, consumers may learn about it over time through usage. More specifically, I consider learning via bad news signals - when a breach occurs, the users of the website detect it with an exogenous probability λ. The detection of a data breach is assumed to be private information of a consumer[1]. This implies that only users of the website may learn about a breach. One potential interpretation of λ is as the probability that a consumer notices fraudulent transactions on his bank account.

The timing of the game is as follows:

- $t = 0$: The website decides the amount, $c(q)$, to invest in data security. The state of security realises.
- $t = 1$: Consumers choose whether or not to use the website given their valuation for the product and the reputation of the firm. A data breach may occur during the period; if it occurs, the users learn about (or detect) it with probability λ. The users revise their belief about the firm's security.
- $t = 2$: Consumers make their usage decisions for the second period. As with the first period, a data breach may occur.

[1] The assumption that learning is private is not crucial to the analysis and is made mainly for computational simplicity. One may also assume that learning is public; i.e. both users and non-users may learn about the data breach. This would correspond to the case of public breach announcements or coverage by the media. The reputation effect of a breach would be larger with public signals. An illustration of the public signal case can be found in policy analysis section, where I discuss the impact of mandatory media notice.

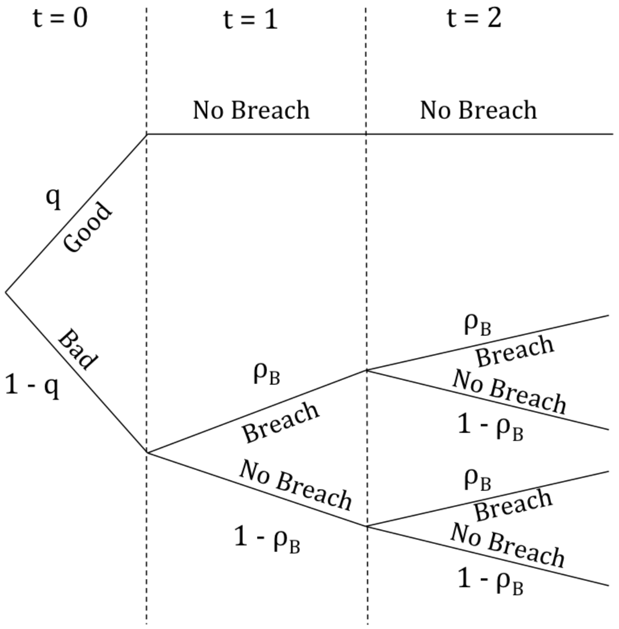

Figure 1: Security investment outcomes

In each period, a consumer has to decide whether or not to use the website. The consumer values the product at v but risks to suffer a financial loss of L in the event of a breach. The consumer's expected utility from using the website within a period is

$$E(U) = v - E(\rho)L$$

Let $U_{1,M}$ (the subscript "M" denotes consumer myopia; this will be discussed later in this section) and U_2 be the consumer's utility from using the website in period 1 and period 2 respectively. The consumer has the following initial belief about the website's investment outcome:

$$Pr(\rho = 0) = q^c$$

Given his belief, the consumer's expected utility from using the website in period 1 is

$$E(U_{1,M}) = v - (1 - q^c)\rho_B L, \tag{1}$$

where $(1 - q^c)\rho_B$ is the expected probability of a breach. The consumer's expected utility in the second period depends on the updated belief of the consumer:

$$Pr(\rho = 0|\text{Breach detected}) = 0, \quad \text{and}$$
$$Pr(\rho = 0|\text{No breach detected}) = \frac{q^c}{1 - \lambda(1 - q^c)\rho_B}.$$

The event where no breach was detected comprises two possible scenarios: (i) there was indeed no breach and (ii) there was a breach but it went undetected. The corresponding expected utilities of engaging with the website in the second period are:

$$E(U_2|\text{Breach detected}) = v - \rho_B L \tag{2}$$

and

$$E(U_2|\text{No breach detected}) = v - \frac{(1-q^c)(1-\lambda\rho_B)}{1-\lambda(1-q^c)\rho_B}\rho_B L. \tag{3}$$

Having set up the expected utilities, we may now proceed to examine the consumer's usage decision. Let us first consider the consumer's problem in period 2. Conditional on learning that there was a breach in period 1, the consumer will choose to continue using the website if his expected utility from usage exceeds the value of his outside option. I assume here that the value of his outside option is 0. Therefore, the consumer decides to use in the website if his valuation for the website lies above

$$\hat{v}_D = \rho_B L. \tag{4}$$

Similarly, in the event that no breach was detected, the consumer uses the website if his valuation exceeds

$$\hat{v}_{ND}(q^c) = \frac{(1-q^c)(1-\lambda\rho_B)}{1-\lambda(1-q^c)\rho_B}\rho_B L. \tag{5}$$

We will now examine the consumer's problem in period 1. If the consumer is myopic, i.e. he only considers his expected utility of usage in the current period, he will choose to use the website if his valuation for the website's product is greater than

$$\hat{v}_M(q^c) = (1-q^c)\rho_B L. \tag{6}$$

With forward-looking or non-myopic consumers, the decision problem in period 1 becomes slightly more complex. Recall that only consumers who choose to use the website have the possibility of learning about the website's data security level (breach detection is a private event). While the possibility of learning does not affect consumers with very high or very low valuations - those with valuation above $\hat{v}_D$ will always use the website and those with valuation below $\hat{v}_{ND}$ will never use the website regardless of whether or not they learn - consumers with valuations lying between $\hat{v}_{ND}$ and $\hat{v}_M$ face a trade-off. Given his initial belief, such a consumer obtains a negative level of utility from using the website in the first period; however, he derives a positive level of utility from usage in the second period in the event where no breach was detected. The forward-looking consumer understands that staying out in the first period precludes participation in the second period, since he learns nothing and therefore will not update his belief. Thus, by not using the website in the first period, he forgoes any positive expected utility that he may obtain in the second period. Taking into account the value of the possibility of learning, the non-myopic consumer's expected utility of using the website in period 1 is

$$E(U_{1,NM}) = E(U_{1,M}) + \delta(1-\lambda(1-q^c)\rho_B)E(U_2|\text{No breach detected}). \tag{7}$$

The forward-looking consumer chooses to engage with the website in the first period if his valuation lies above

$$\hat{v}(q^c) = \frac{(1-q^c)(1+\delta(1-\lambda\rho_B))}{1+\delta(1-\lambda(1-q^c)\rho_B))}\rho_B L, \tag{8}$$

where δ is the discount factor. Notice that $\hat{v} \geq \hat{v}_M$ as one may expect; a non-myopic consumer may choose to use the website in the first period (albeit obtaining negative expected utility) in order to have the opportunity to learn and, hence, the option to use the website in the second period. The consumers' valuation thresholds are illustrated in Figure 2. Notice that although the usage threshold in the second period when no breach was detected is $\hat{v}_{ND}(q^c)$, the fraction of users in the market is given by $1-\hat{v}(q^c)$. The consumers with valuation between $\hat{v}(q^c)$ and $\hat{v}_{ND}(q^c)$ would like to participate in the second period when no breach was detected; however, as they did not use the website in the first period, they do not learn anything about the website's level of security and continue not to participate.

Implicit in the above analysis is the assumption that a consumer who used the website in the first period would not incur any loss in the second period if he decides to stop using the website. This would be true when the website deletes the consumer's information once he quits the site or when the information collected in the first period becomes obsolete in the second (for example, the consumer's payment card information would no longer be valid if he had his card replaced). In reality, websites often continue to retain the consumer information they have collected even after a consumer has stopped using their sites. The model can be easily modified to depict this scenario. This can be done by introducing an additional term to capture the loss that a user in the first period may expect to incur in the second period even after leaving the website. The nature of the consumer's problem would remain unchanged, as long as the expected loss from not using is smaller than that of using the website in the second period. In this case, for any given initial level of reputation, there will be fewer users in the first period but more users in the second period following breach detection. For the rest of this paper, I will focus on the case where the consumer would not incur a loss in the second period if he is not using the website.

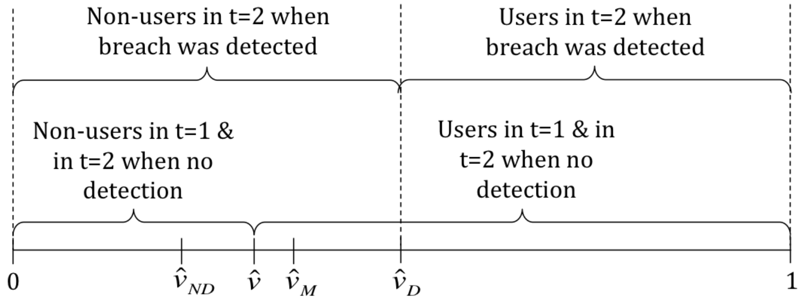
Non-users in t=2 when breach was detected
Users in t=2 when breach was detected
Non-users in t=1 & in t=2 when no detection
Users in t=1 & in t=2 when no detection
0
$\hat{v}_{ND}$
$\hat{v}$
$\hat{v}_M$
$\hat{v}_D$
1

Figure 2: Valuation thresholds

We now turn to the website's decision problem. Given the consumers' usage thresholds, the website needs to determine the amount of security investment to make, i.e.

$$\max_{q^f} \pi(q^f, q^c),$$

where the website's profit function is by

$$\begin{aligned}\pi(q^f, q^c) = (1 - \hat{v}(q^c))\, r + \delta \big[\lambda(1 - q^f)\rho_B \,(1 - \hat{v}_D) \\ +(1 - \lambda(1 - q^f)\rho_B)\,(1 - \hat{v}(q^c)) \big]\, r - c(q^f).\end{aligned} \tag{9}$$

This gives the following profit-maximising condition

$$c'(q^*) = \delta\lambda\rho_B \,(\hat{v}_D - \hat{v}(q^c))\, r \tag{10}$$

3 Equilibrium Analysis

I analyse the Bayes-Nash equilibria of this game. The equilibria of this game are pinned down by the consumers' usage thresholds, $\hat{v}, \hat{v}_D$ and $\hat{v}_{ND}$, and the website's profit maximising condition. We also require the consumers' belief to be consistent, i.e. $q^c = q^*$ [2]. Plugging the consumers' usage thresholds into the website's profit-maximisation condition and imposing consistent beliefs, we obtain the following condition

$$c'(q^*) = \delta\lambda\rho_B \left(\frac{(1+\delta)q^*}{1 + \delta(1 - \lambda(1 - q^*)\rho_B)} \right) \rho_B L r \tag{11}$$

From here on, I will refer to the right hand side of Equation (11) as the marginal benefit of investment. The condition basically states that a profit-maximising firm should invest up to the point where the marginal cost of investment equals to the marginal benefit. So long as the marginal cost of not investing is zero (i.e. $c'(0) = 0$), this condition is trivially satisfied at $q^* = 0$.

Lemma 1. (No Investment Equilibrium).
There always exists an equilibrium where the website does not invest in data security, i.e. $q^* = 0$.

The intuition behind this lemma is simple. When the website does not invest in data security, the resulting probability of data breach is ρ_B with certainty. Since consumers have rational expectations, they believe that the probability of a breach is always ρ_B and there is no updating of beliefs after the first period. The valuation thresholds - which are the same in both periods since consumers' beliefs are constant - maximise the consumers' utility given their beliefs. It is also easy to verify that the website is playing best-response to the consumers' strategies and beliefs. In general, taking the consumers' beliefs as given, the benefit of investing in data security arises from the reduction in the likelihood that a breach occurs and is detected by the consumers. This reduction, however, only benefits the website to the extent that the consumers' valuation threshold is smaller under the case of where no breach is detected than that where a breach is detected. When consumers believe that the firm is not investing with

[2] In the appendix, I also analyse the case where the Bayesian consumers have exogenous beliefs.

certainty, the valuation thresholds are the same whether or not a breach is detected. Thus, investing in data security only results in higher cost for the website without increasing its profit and it is optimal for the firm not to invest in data security.

There may also be other equilibria in this model. The existence of and the equilibrium level of investment in these equilibria depend on the shape of the marginal cost and marginal benefit functions. In essence, the website is willing to invest in data security when the marginal benefit of doing so outweighs the marginal cost. The benefit of investment arises from the reduction in expected consumer turnover. When consumers learn that a breach has occurred, they update their beliefs about the website's security unfavourably (i.e. a decline in the website's reputation for quality). This leads to consumer turnover, and a loss of revenue in the next period. As evident from equation (11), this reputation cost (or the marginal benefit of investment) is positive for all values of q, so long as the user-generated revenue, r, is positive. Further, one can verify it is concave in q for all $q \in [0, 1]$, which implies that the increase marginal benefit of investing tapers off as the level of investment increases.

Now, consider an increasing and convex cost function, $c(.)$, with the usual properties:

1. $c(0) = c'(0) = c''(0) = 0$;
2. $c(1) = c'(1) = +\infty$;
3. $c'(q) \geq 0$ and $c''(q) \geq 0$ for all $q \in [0, 1]$.

The proposed cost structure, though simple, encapsulates several aspects of the data security investment problem. The convexity of the cost function reflects how improving the level of security becomes increasingly difficult (and costly) as the firm becomes more secure. Indeed, a firm that is completely unsecured can easily raise its level of security by introducing measures that are relatively costless (such as stronger passwords and email encryption). By contrast, a firm that already possesses a strong security posture may have to purchase more costly data protection softwares or engage consultants to further improve its level of security. Another aspect of data security investment - perfect security is not optimal - is captured by the second property of the cost function. It is too costly, if not altogether impossible, for a firm to reduce the probability of a data breach to 0. A similar assumption that perfect security cannot be attained with a finite amount of investment is also described in Gordon and Loeb (2002). Therefore, given the cost structure, the profit-maximising level of investment, q^*, should lie between 0 and 1. The following proposition sums up the above discussion.

Proposition 1. (Positive Investment Equilibrium).
Suppose the cost function satisfies the properties as stated above[3] There exists a stable

[3] In the case where $c''(0) = 0$ (which violates property 1), the positive investment equilibrium only exists if r exceeds a certain threshold, $\hat{r}$. This threshold $\hat{r}$ can be shown to be equal to $c''(0)\left(\frac{1 + \delta(1 - \lambda\rho_B)}{(\delta\lambda\rho_B^2)(1 + \delta)L}\right)$.

equilibrium with a positive level of investment for all values of r. The equilibrium level of investment, q^*, satisfies equation (11).

At this point, one might wonder what are the factors that affect the level of investment in equilibrium. Observe from (11), for a given marginal cost function and level of revenue r, the level of investment is likely to be higher for larger values of the consumer data breach losses (l), the discount factor (δ) and the probability of breach (ρ_B). These parameters affect the cost of a data breach to a firm. An increase in L magnifies the fraction of consumers that the firm loses ($\hat{v} - \hat{v}_D$) in the case of breach detection; a higher δ indicates that the firm weighs the potential loss of future business more; a larger ρ_B signifies that a data breach is more likely. The effect of a change in λ on the equilibrium level of investment is less certain. The parameter λ can be thought of as a measure of consumer learning; the higher the λ, the more probable that a consumer would detect a data breach. The impact of an increase in λ on the equilibrium will be discussed in greater detail in the policy analysis section.

I conclude this section with a few welfare results. Let us first look at how the two equilibria compare in terms of social welfare. Setting the gains of cybercriminals aside, social welfare can be decomposed into the sum of consumer surplus and the website's profit. Consumer surplus is always increasing with the level of security. First, more consumers use the website (in both periods) when the it invests in data security. Second, the expected utility of the users is higher, since investment in security lowers the expected breach losses that the consumers face. Hence, consumer surplus is higher in the positive investment equilibrium. One can also verify that the website's profit is higher in this equilibrium. For any given level of investment made by the website, its profit is an increasing function of consumers' belief, q^c. When the consumers believe that q^c is high (or the expected probability of breach is low), more consumers become users of the website; this generates a higher level of revenue for the website. Consider now the no investment equilibrium where $q^f = q^c = 0$. The website's profit would be higher in the case where it does not invest ($q^f = 0$) but the consumers (mistakenly) believe it to be investing at a level $q^c = q^*_+$; that is, $\pi(0, q^*_+) > \pi(0, 0)$. Further, by the revealed preference argument, we know that the website's profit is maximised at q^*_+ when the consumers' belief is $q^c = q^*_+$; that is, $\pi(q^*_+, q^*_+) \geq \pi(0, q^*_+)$. Thus, we have established that the website's profit is higher under the positive investment equilibrium. Since both consumer surplus and the website's profit are higher, social welfare is also higher in the positive investment equilibrium. In other words, the positive investment equilibrium Pareto dominates the no investment equilibrium.

Next, let us examine how the level of investment at the market equilibrium compares with that at the social optimum. Given that the website does not internalise the loss that data breaches impose on consumers, one may expect the market equilibrium to feature an underinvestment in security relative to the social optimum. Indeed, it can be shown that the social welfare maximising level of investment, q^s, is higher than the market equilibrium level of investment, q^*. More formally, one can verify that the equilibrium profit of the website is increasing in the level of investment at $q = q^*$.

This, coupled with the fact that consumer surplus is always increasing in the level of investment, implies social welfare is also increasing with investment at $q = q^*$. The following proposition summarises main welfare results of the discussion above:

Proposition 2. (Welfare Results).
(i) The positive investment equilibrium Pareto dominates the no investment equilibrium.
(ii) In both equilibria, the website underinvests in security relative to the social optimum; i.e. $q^* < q^s$.

We have established in this section that the positive investment equilibrium is stable and Pareto dominates the zero investment equilibrium. Therefore, I will focus the analysis on the positive investment equilibrium for the remainder of the paper.

4 Policy Analysis: Mandatory Data Breach Notification

In the absence of regulation, the website's incentive to invest in security stems from the potential loss of consumers in the event of a breach. Since the website only loses consumers when they learn about a breach, the extent to which the market is able to instill discipline in the firm is proportional to the probability of data breach detection. One may expect the level of data breach detection by users to be low for various reasons. Consumers may not go through their monthly card statements thoroughly; even when they do, it is often difficult for them to identify where the charges to their cards came from. The "merchant decriptor" on monthly card statements is said to be "frustrating brief" - it is limited to between 26 and 28 characters (The New York Times (2010, August 21)). This makes it hard for even the most meticulous of consumers to distinguish between legitimate and fraudulent transactions. Furthermore, consumers typically interact with multiple firms; they may not be able to ascertain the source of the breach even if they were able to detect one. Thus, one way of inducing the website to invest more could be to introduce policy measures that increase the rate of breach detection or learning by consumers.

One such policy measure is mandatory breach notification. Data breach notification laws are present in many states in the United States; however, their requirements differ from one state to another. The increase in the frequency and scale of data breaches in the recent times has prompted extensive discussions over the enactment of a data protection and breach notification legislation at the federal level. Several bills have been proposed to date, though none has yet to be enacted. An example is the Data Security and Breach Notification Act of 2015 (S. 177). The proposed bill legally obliges firms to provide timely notice both to the FTC and to consumers in the event of a breach. The notification can take the form of a written correspondence, a telephone notice or an email notice. Should the data breach affect more than 5,000 individuals in a given State, the breached firm is also obliged to provide media notice in that State. Under the stipulation of this act, deliberate concealment of data breaches would result in criminal penalties (fine or imprisonment or both).

In the section, I examine the impact of the proposed breach notification law on the level of security and social welfare in equilibrium. To enrich the analysis, I also introduce the possibility for consumers to self-insure (i.e. they may take action to reduce the magnitude of their losses) upon learning that a breach has occured. Romanosky et al. (2011) give several examples of precautionary measures consumers can take upon learning that their information has been comprised. These measures include informing their banks, closing unused accounts, and applying for credit freezes and fraud alerts. Let α denote the fraction of losses that may be avoided by the consumer and l denote the amount of losses when he takes no action. When the consumer can self-insure, the amount of losses that the he suffers as a result of a breach is thus $L = (1 - \lambda\alpha)l$. The possibility of self-insurance creates a link between breach detection and the magnitude of losses, and leads to more nuanced policy implications. This will be discussed in greater detail in the following subsections.

4.0.1 Individual Notice

Let us first analyse the case where the website is required to provide individual notice only; i.e., for data breaches involving fewer than 5,000 records. In the context of my proposed framework, mandatory breach notification translates into setting the probability of breach detection, λ, to 1. This increase in λ generates two effects: a reputation effect and a crowding out effect. The crowding out effect, whose strength depends on the fraction of avoidable loss (as captured by the variable α), countervails the learning effect. The overall impact of the regulation on the positive investment equilibrium hence depends on the relative magnitude of the two effects. The following proposition states the conditions under which a mandatory notification law would raise the level of investment in equilibrium.

Proposition 3. (Level of Investment with Individual Notice).
Let q^*_{in} denote the level of investment at the positive investment equilibrium. $q^*_{in} \geq q^*$ if one of the following conditions are satisfied:
(i) $0 \leq \alpha < \dfrac{1}{1+\tilde{\lambda}}$, or
(ii) $\dfrac{1}{1+\tilde{\lambda}} \leq \alpha < \dfrac{(1+\delta)}{(1+\delta+\tilde{\lambda}(1+\delta(1-\rho_B))}$ & $r \geq c'(q^*)\left(\dfrac{1+\delta(1-(1-q^*)\rho_B}{\delta\rho_B^2(1+\delta)q^*(1-\alpha)l}\right)$.
Otherwise, $q^*_{in} < q^*$.

To facilitate the discussion of this result, let us first shut down the crowding out effect by supposing that $\alpha = 0$. This corresponds to the case where the consumers are unable to avoid any losses that result from a data breach. When this is so, the data breach notification law raises the level of investment in equilibrium. The main intuition underlying this result is as follows: mandatory individual notice increases the probability of breach detection, λ, from $\tilde{\lambda}$ to 1, thereby alleviating part of the informational asymmetry between the website and the consumers. One can think of λ as a measure of consumer learning - the occurrence of a breach conveys information about the level of security at the website and λ determines the extent to which this information is transmitted to the consumers. When λ is high, consumers are more likely to detect a breach and consequently the website is more likely to lose consumers when a breach occurs. The increase in λ also lowers the first period valuation threshold, $\hat{v}$ of the consumers. Recall that consumers incorporate the option value of learning when deciding whether or not to use the website in the first period. This value of learning is higher when λ is large. Thus, by raising λ, the breach notification law makes the option of learning more valuable to consumers and increases the level of participation in the first period. This widens the gap between $\hat{v}$ and $\hat{v}_D$, which means that the website has more to lose in the event of a breach. To sum up, the increase in consumer learning strengthens the firm's incentive to invest via two channels. First, it raises the probability that the firm loses consumers when a breach occurs. Second, it increases the amount at stake for the website (as $\hat{v}$ is now higher).

Whereas the breach notification law unambiguously increases the firm's incentive to invest when $\alpha = 0$, its impact is less certain for positive values of α. α is the fraction of breach-related losses that a consumer can take action to avoid, conditional on learning about the breach; the eventual amount of losses that the consumer suffer in the event of a breach is given by $(1 - \lambda\alpha)l$. That said, observe that a larger value of λ also implies a smaller eventual amount of losses suffered by the consumer. One can think about $\lambda\alpha$ as a measure of the consumer's ability to insure himself against breach-related losses. In raising λ, mandatory breach notification strengthens the consumer's ability to self-insure. Indeed, there is some empirical evidence that data breach notification laws may reduce the occurence of identity thefts, which one could interpret as an indication of better consumer self-insurance[4] Because better self-insurance leads to lower expected losses, more consumers are now willing to participate at any given belief; i.e. $\hat{v}$ and $\hat{v}_D$ are both lower. The reduction in breach losses has a stronger effect on $\hat{v}_D$ relative to $\hat{v}$, however. This is because the expected probability of a breach is higher when a breach has been detected. Since the reduction in breach losses only matter if a breach were to occur, it induces a larger decrease in $\hat{v}_D$ as compared to $\hat{v}$. This dampens the website's incentive to invest in security as it now has less to lose when a breach occurs. Therefore, one can say that an increase in the consumer's ability to self-insure crowds out the website's investment in data security.

In a nutshell, the mandatory breach notification requirement generates two effects - a learning effect and a crowding out effect. The learning effect causes $\hat{v}$ to fall relative to $\hat{v}_D$, hence strengthening the firm's incentive to invest; the crowding out effect causes $\hat{v}_D$ to decrease more than $\hat{v}$, weakening investment incentives. The overall impact of

[4] Using state-level variations in the adoption of data breach notification laws in the United States, Romanosky et al. (2011) estimated the impact of the passage of these laws on identity theft. They found that the adoption of data breach notification laws reduced the occurrence of identity thefts by 6.1%.

the breach notification law depends on the relative strength of these effects. As can be seen from the propositions, the learning effect dominates for small values of α while the reverse is true for large values of α. This is due to the fact that the crowding out effect is stronger for large values of α, since an increase in α implies that consumers can better insure themselves against breach losses. For intermediate values of α, the impact of the breach notification law on investment incentives may also depend on another factor - the level of user-generated revenue, r. When r is large, the value of a consumer to the website is higher; correspondingly, a data breach is more costly for the firm. This strengthens the learning or reputation effect of the breach notification law and induces the website to invest more in security. Hence, when r is sufficiently high, the learning effect continues to dominate for intermediate values of α and the equilibrium level of investment is higher relative to when there is no regulation. The graphical representation of the various cases discussed in Proposition 4 can be found in figures 3 to 6.

4.0.2 Media Notice

In the analysis thus far, I have assumed consumer learning to be private - only the consumers who have been affected by the breach can learn about it. Consumers who chose not to use the website cannot be affected by a breach and thus they do not learn anything about the website's level of security. This assumption continues to hold with mandatory individual breach notification. With media notice, however, data breaches become publicly known - both users and non-users of the website learn about them.

In addition to its impact on the website's incentive to invest, media notice also alters the consumers' usage behaviour. Recall the discussion about myopic versus forward-looking consumers. The forward-looking consumer foresees that he would like to use the website in the second period if no breach was detected (which would lead to a favourable update in his belief) in the first period; however, he knows that he would not learn if he does not participate in the first period. Due to the positive value of learning, he may still choose to use the website even when his expected within-period utility for the first period is negative. When data breach notification is made through the media, the learning motive for participating in the first period is eliminated. The first period valuation threshold now corresponds to that of the myopic consumer as in equation (6) (with $\lambda = 1$). The period 2 thresholds remain unchanged (again with λ set to 1).

Although the consumers' valuation thresholds in the second period is the same as before, the relevant market size for the website (in the case where no breach was detected) has changed. The fraction of consumers who chooses to use the website's services in this case is now given by $1 - \hat{v}_{ND}(q^c)$. As information on data breaches is now public, consumers who did not participate in the first period are also able to learn about the website's state of security. Some of these consumers - more specifically, those with valuations that lie between $\hat{v}(q^c)$ and $\hat{v}_{ND}(q^c)$ - will decide to use the website in the second period. The equilibrium when the firm is required to provide media notice

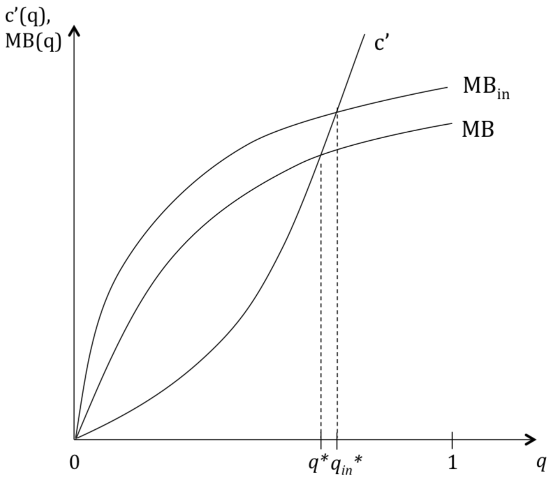

Figure 3: Small α

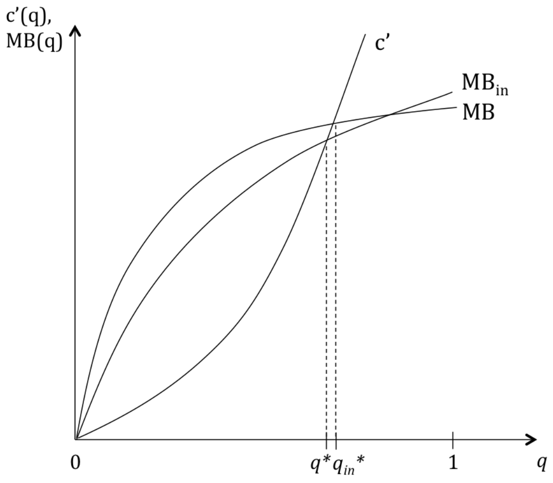

Figure 4: Intermediate α, large r

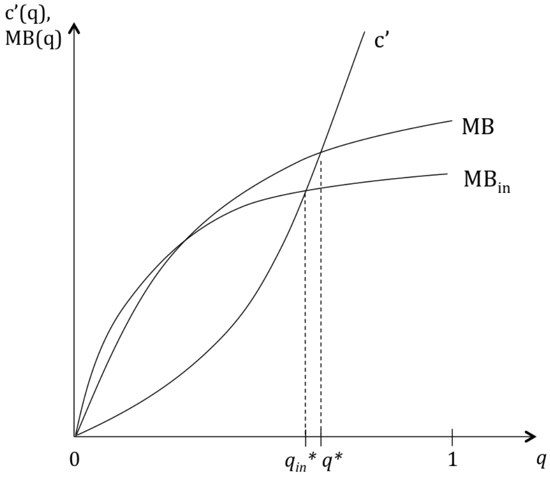

Figure 5: Intermediate α, small r

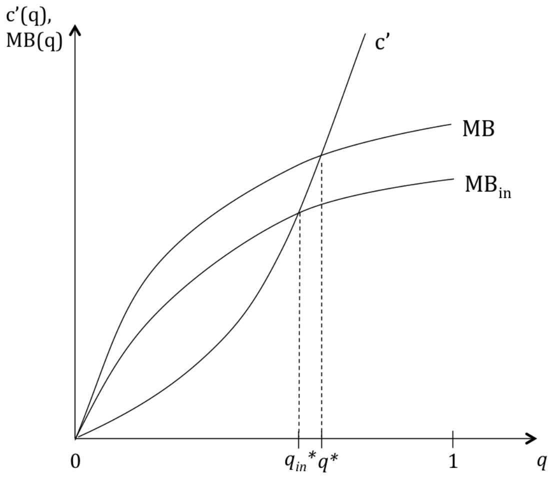

Figure 6: Large α

is pinned down by the following condition:

$$c'(q^*_{mn}) = \delta\rho_B\left(\hat{v}_D - \hat{v}_{ND}(q^*_{mn})\right)r, \tag{12}$$

where

$$\hat{v}_{ND}(q^*_{mn}) = \frac{(1-q^*_{mn})(1-\rho_B)}{1-(1-q^*_{mn})\rho_B}\rho_B(1-\alpha)l \quad \text{and}$$
$$\hat{v}_D = \rho_B(1-\alpha)l.$$

Stipulating media notice boosts the website's incentive to invest further, as compared to the case where only the provision of individual notice is mandated. As previously mentioned, media notice enables non-participants to learn about the firm's level of security; some of these consumers may choose to participate in the second period if no breach was announced during the first period. Consequently, the website has more to gain from reducing the probability of experiencing a data breach under media notice. In other words, the marginal benefit of investing in data security has increased. In fact, it can be shown that the marginal benefit of investment with media notice lies everywhere above that when only individual notice is required. This implies a relatively higher level of investment in equilibrium. The following proposition summarises the main results of the discussion.

Proposition 4. (Media Notice).
Let q^*_{mn} denote the equilibrium level of investment when media notice is required in addition to individual notice. By stipulating mandatory media notice,
1. $q^*_{mn} \geq q^*$ if either of the following conditions are satisfied:

(i) $0 \leq \alpha < \dfrac{1}{1+\tilde{\lambda}}$, or

(ii) $\dfrac{1}{1+\tilde{\lambda}} \leq \alpha < \dfrac{(1+\delta)(1-\tilde{\lambda})+\tilde{\lambda}\rho_B}{(1+\delta)(1-\tilde{\lambda}^2)-\delta\tilde{\lambda}\rho_B(1-\tilde{\lambda})+\tilde{\lambda}^2\rho_B}$
$\& \; r \geq c'(q^*)\left(\dfrac{1-(1-q^*)\rho_B}{\delta\rho_B^2 q^*(1-\alpha)l}\right)$.

Otherwise $q^*_{mn} < q^*$.
2. $q^*_{mn} \geq q^*_{in}$ for all parameter values.

Welfare Implications

In the preceding discussion, we have focused on the impact of a mandatory breach notification law on the website's investment incentives. While the effect of the regulation on the level of investment is an important consideration, a more complete evaluation of the policy would require an examination of its implication on social welfare. The two concepts are in fact closely related. We have established that the website under-invests in equilibrium as compared to the social optimum. Hence, one might expect social welfare to be higher whenever the notification law raises the level of investment in equilibrium. The following proposition formalises this intuition.

Proposition 5. (Welfare Implications of Mandatory Notification)
Under the proposed mandatory breach notification act, the level of social welfare is
(i) higher as compared to the unregulated market outcome whenever the regulation brings about an increase in the equilibrium level of investment;
(ii) always higher with the additional requirement of media notification relative to when only mandatory individual notice is imposed.
The effect of the proposed regulation on social welfare is ambiguous whenever it reduces the level of investment in equilibrium.

Mandatory breach notification affects both consumer surplus and the website's profit. The regulation effect on consumer surplus can be decomposed into a direct and an indirect component. Holding the level of investment constant, the direct impact of mandatory breach notification on consumer surplus is positive. First, the proposed regulation reduces the losses that consumers face in the event of a breach. Since the website is obliged to inform its customers of data breaches, consumers are able to take actions to mitigate part of the potential losses; consequently, the level of consumer surplus is increased. Second, mandatory breach notification can help to eliminate excessive usage in the second period, thereby raising the level of consumer surplus. Following a data breach in the first period and in the absence of mandatory breach notification, a fraction of the consumers (whose valuation lie between $\hat{v}$ and $\hat{v}_D$) derive negative utility from using the website in the second period; however, they may continue to participate because they failed to detect the breach in the first period. In other words, without mandatory breach notification, there may be socially excessive participation in the second period. By ensuring that consumers learn about data breaches, this excessive level of participation can be eliminated.

The indirect effect of mandatory breach notice on consumer surplus may be either positive or negative, and occurs through its impact on the equilibrium level of investment. Consumer surplus is increasing in the level of investment - a higher level of investment results in smaller expected losses from data breaches and in a higher level of participation. We know from propositions 4 and 5, however, that mandatory breach notification may not necessarily boost the level of investment in equilibrium. The sign of the indirect effect is thus ambiguous and depends on the regulation's impact on the investment incentive of the website. Since the direct effect of the policy on consumer surplus is always positive, the level of consumer surplus is higher whenever mandatory notification raises the level of investment in equilibrium. When the regulation reduces investment incentive, however, its impact on consumer surplus would depend on the relative magnitudes of the direct and indirect effects.

The impact of mandatory breach notification on the firm's profit can also be decomposed, with a slight abuse of terminology, into a "direct" and an indirect effect. The "direct" effect here refers to the regulation's impact on profit, holding the level of investment constant. The "direct" effect is comprised of two sub-components. On the one hand, the rise in probability of breach detection increases the reputation cost of a breach, and hence reduces the website's profit. On the other hand, the increase in λ brings about a higher level of participation in the market (by lowering the valuation thresholds), creating a positive impact on profit. Overall, the "direct" effect can be shown to be positive. The indirect effect of the regulation on profit occurs via the policy's impact on the level of consumer participation. The level of participation is increasing in both the probability of breach detection, λ, and the level of investment in equilibrium. By raising λ, mandatory breach notification reduces the expected breach-related losses that consumers face, which augments the level of consumer participation. The regulation may also reduce the amount of investment in equilibrium, however, thereby lowering the level of consumer participation. Consequently, the indirect effect of the regulation on the firm's profit is ambiguous in the case where the equilibrium level of investment is decreased.

It can further be shown that imposing media notice on top of mandatory individual notice is welfare enhancing; the level of social welfare is always higher when the website is obliged to provide both individual and media notice as compared to when it is required to give individual notice only. Intuitively, this may be because media notice eradicates another layer of information asymmetry - the one between the website and non-users - alleviating the underinvestment problem present in the market.

Before concluding this subsection, one remark may be worth mentioning. In the analysis, we have interpreted α as the fraction of avoidable breach-related losses. Implicit in this interpretation is that these losses has not yet been generated. It is also possible - and is the case for losses due to payment card fraud - that these losses have already been incurred by the consumers but may be passed on by the consumers to another party. Indeed, conditional on detecting fraudulent transactions, consumers are typically insured against the resulting losses by their banks. The welfare impact of the breach notification law is less favourable when α partially or fully reflect the fraction of breach-related losses that is transferred from the consumers to another part of the society. Recall from the earlier discussion that the enactment of the mandatory breach notification generates two effects: a reputation or learning effect and a crowding out effect. The crowding out effect, which dampens the firm's incentive to invest, is stronger for higher values of α. As such, the regulation is less likely to be welfare-enhancing under these circumstances. If, in addition, the larger value of α merely reflects a larger portion of losses that is borne by the banks, the welfare implication of the policy would be even less favourable. In this case, the increase in λ due to the introduction of the law only results in consumers transferring a larger fraction of their losses to the banks and leaves the total amount of losses suffered by the society as a whole unchanged. A better allocation of breach liability between the website, the consumers and the banks may be necessary in order for mandatory breach notification to be welfare improving.

5 Strategic Interactions in Security Investments

The first part of this paper has focused on the interaction between the website and the consumers, illustrating how the reputation effect drives the website's investment in security. In addition to the reputation effect, interdependencies in security investments may also affect the website's incentive (and hence the overall level of security). In this section, I introduce security interdependencies into the framework by considering the role of the consumers' bank in data breaches. I explore how these interdependencies affect the investment incentives of both the website and the bank and their implications on the overall level of security.

Data breaches at firms reportedly account for more stolen money from banks than robberies (Gorman and Evan (2009)). Financial information compromised in data breaches is frequently used to counterfeit payment cards, which are then used for cash withdrawals at ATM or for payment in physical and/or internet transactions (Sullivan (2010)). In many cases, the banks bear most, if not all, of the liability for the fraudulent payments charged to these cards[5]. The American Bankers Association attributes the rise in fraud losses in 2014 to the large scale retail data breaches (American Bankers Association (2016, January 2016)). In addition to the stolen money and fraudulent charges, banks may also incur the costs of card re-issuance and for the operation of call centres to handle consumer enquiries

While banks may be regarded as victims of data breaches, it is important not to overlook their role in determining the outcome of these breaches. The incidence of fraud following a data breach also depends on the strength of the banks' security measures. Recent data breaches in the U.S. highlight how the inadequacy of its legacy payment system has facilitated card fraud. The use of the magnetic stripes to store information on payment cards, for example, makes it easy for criminals to counterfeit these cards (Kerber (2013)). The relatively weak payment approval process also contributes to the higher rates of fraud in the U.S. as compared to other countries (Sullivan (2010)).

The above discussion demonstrates how the overall level of security may be jointly determined by the website's and the banks' investment levels, implying the presence of interdependencies in security investments. These interdependencies generate strategic interactions between the parties' investment strategies and may have important implications for policy making. In the subsections that follow, I extend a simplified version of the baseline model to include the consumers' bank. I then characterise the equilibria of the investment game for two different types of bank security measures and study the implications of a minimum bank security standard and the mandatory breach notification law.

5.1 Extended model: Bank as a Third Player

Consider a simplified version of the baseline model with a representative consumer and a website and introduce into this framework the consumer's bank. Let us assume that the relationship between the consumer and his bank is pre-existing (i.e. I will not model the consumer decision to use the bank). The bank provides the consumer with a payment card, which he uses for transactions with the website. When used, the card's information is collected and stored by the website. In order not to complicate the analysis, I will assume that data breaches only occur at the firm but not at the bank. The consumer expects his payment card information to be breached with probability $(1 - q^c)\rho_B$, where q^c denotes his initial belief that the website is secure ($\rho = 0$).

Like the website, the bank may invest in security. The bank's investment determines the likelihood that fraud may be committed with stolen payment card information. Let $1 - \gamma$ be the probability that a data breach at the website results in payment

[5] Depending on the type of card and the type of transaction, the banks are typically liable for the fraudulent charges incurred, provided that consumers detect and report the losses before a certain deadline. For more information on payment card liability, visit: http://www.consumer.ftc.gov/articles/0213-lost-or-stolen-credit-atm-and-debit-cards.

card fraud, where γ reflects the bank's level of security. γ is assumed to be observed by the consumer. This is on the grounds that a consumer can usually obtain a good gauge of his bank's security level from the payment-related security measures (e.g. the requirement of PINs or passwords for transactions and the blocking of unusual transactions) adopted by the bank. Taking into account the bank's security level, the consumer expects to incur a loss with probability $(1 - q^c)\rho_B(1 - \gamma)$ if he chooses to use the website in the first period.

I distinguish between two types of security measures that may be undertaken by the bank. The first type of measures involves the active screening of transactions. The requirement of a one-time PIN in addition to credit card credentials for transactions, for example, is an "active" security measure. With "active" measures, the consumer becomes aware of fraud attempts made on his card; this represents an additional channel through which the consumer may learn about a data breach at the website. The second type of measures (passively) lowers the likelihood that a fraud may be committed. An example of "passive" measures is the adoption of "chip-and-PIN" payment cards. "Chip-and-PIN" cards are more difficult to counterfeit, thereby lowering the incidence of fraud (Sullivan (2010)). With "passive" measures, the consumer does not learn about a data breach when it does not result in a loss. Figures 7 and 8 provide graphical illustrations of how the game proceeds in first period for "passive" and "active" measures correspondingly.

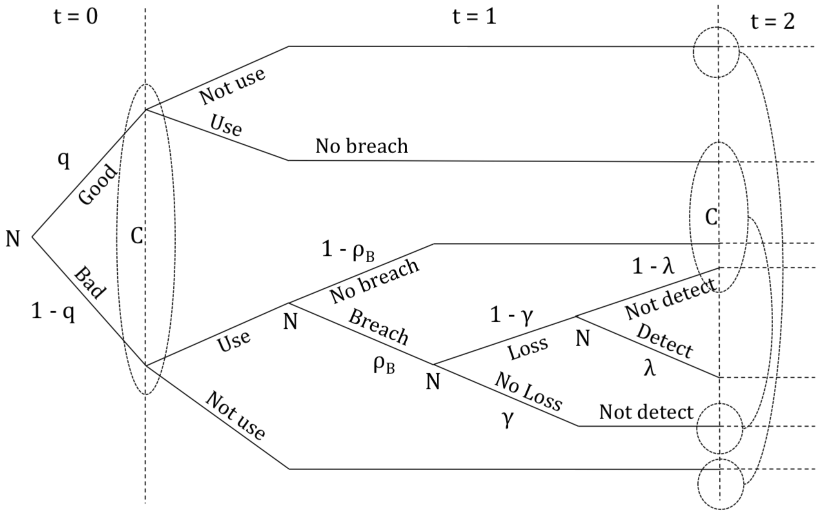

Figure 7: "Passive" measures

Timing

More formally, the timing of the new game is as follows:

- $t = 0$: The website and the bank simultaneously decide on how much to invest in data security.

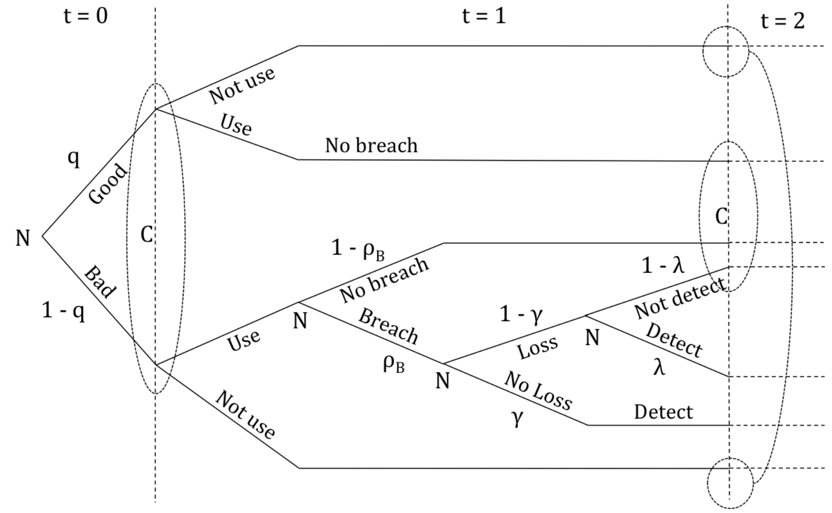

Figure 8: "Active" measures

- $t = 1$: The consumer observes the bank's level of security γ and forms rational expectations over website's level of security. During the period, a data breach may occur at the website and it may generate a loss for the consumer. If a breach occurs, the consumer may either learn about it via his bank (if "active" security measures are in place) or via his own detection of a loss. The consumer revises his belief about the state of security at the website.

- $t = 2$: The consumer makes his usage decision for the second period. As in first period, the consumer may suffer a loss if a breach occurs.

5.1.1 Strategies

I begin with the consumer's problem. The consumer faces the same trade-off as before - he derives a positive level of utility from using the website but exposes himself to potential data breach losses in doing so. Like in the case with consumer self-insurance, the loss that that consumer incur when a breach occur is given by $(1 - \lambda\alpha)l$. Instead of interpreting α as the fraction of loss that the consumer can mitigate, let α represent the fraction of the loss for which the bank is liable. In other words, the bank partially insures the consumer against fraud losses. Again, the consumer is only insured when he detects the loss (and reports it to his bank). The utility that the consumer obtains from using the website in each period is

$$E(U) = v - E(\rho)(1 - \gamma)(1 - \lambda\alpha)l.$$

Observe that the probability that the consumer incurs a loss, given by $E(\rho)(1 - \gamma)$, now depends on both the website's and the bank's level of investment.

Assume that the consumer's valuation for the website's product is sufficiently high such that he would participate in both periods whenever no loss is detected in the first period; that is

$$v \geq (1-q^c)\rho_B(1-\gamma)(1-\lambda\alpha)l,$$

Notice that the right-hand side of the expression corresponds to the expected loss from using the website in the first period. Since the consumer updates his belief favourably when no loss is detected, it follows that his expected utility (given his updated beliefs) of participating in the second period would always be positive if the above condition is satisfied.

There may also be a case where the consumer continues to use the website, despite detecting a breach in the first period. This occurs when his bank's level of security is sufficiently high; that is when γ exceeds

$$\overline{\gamma} = 1 - \frac{v}{\rho_B(1-\lambda\alpha)l}. \tag{13}$$

This suggests that a moral hazard problem may exist between the bank and the consumer - by providing a high level of security, the bank may indirectly encourage consumer to engage in the risky behaviour of using an insecure website.

Consider now the website's security investment problem. The website's profit depends on the nature of the bank's security measures. If the bank's security measures are "passive", the website's profit function is given by

$$\pi_p(q^f,\gamma) = \begin{cases} (1+\delta-\delta\lambda(1-q^f)\rho_B(1-\gamma))r - c(q^f) & \text{if } \gamma \leq \overline{\gamma} \\ (1+\delta)r - c(q^f) & \text{otherwise.} \end{cases}$$

When the bank's security is sufficiently low (i.e. $\gamma \leq \overline{\gamma}$), the introduction of the bank into the model does not alter the fundamental trade-off faced by the website - the website weighs the benefit of a lower level of consumer turnover against the cost of investment. When the bank is highly secure, however, the consumer always remains with the website and the benefit of investment is dissipated. The optimal level of investment satisfies the following condition:

$$c'(q_p^*) = \begin{cases} \delta\lambda\rho_B(1-\gamma)r & \text{if } \gamma \leq \overline{\gamma} \\ 0 & \text{otherwise.} \end{cases}$$

Observe that, for $\gamma \leq \overline{\gamma}$, the website's marginal benefit of investing in security is decreasing with γ. The more secure the bank is, the lower the probability that a breach leads to a loss and the less likely the website loses the consumer as a result of poor data security.

When the bank's security measures are "active", the website's profit function is given by:

$$\pi_a(q^f,\gamma) = \begin{cases} (1+\delta-\delta(1-q^f)\rho_B(\lambda(1-\gamma)+\gamma))r - c(q^f) & \text{if } \gamma \leq \overline{\gamma} \\ (1+\delta)r - c(q^f) & \text{otherwise.} \end{cases}$$

The website's profit function in the case where the bank's security measures are "active" is identical to that when the measures are "passive", except that the probability of losing the consumer when a breach occurs is higher when the measures are "active". This is due to the higher probability of breach detection arising from the bank's monitoring. The consumer learns about a breach with certainty when it does not result in a loss (i.e. when fraudulent transactions are screened out by the bank). The website's optimal level of investment satisfies the following first-order condition:

$$c'(q_a^*) = \begin{cases} \delta\rho_B(\lambda(1-\gamma)+\gamma)r & \text{if } \gamma \leq \overline{\gamma} \\ 0 & \text{otherwise.} \end{cases}$$

Whereas the website's level of investment is decreasing in γ when the bank's security measures are "passive", it is increasing in γ when the measures are "active". Put differently, there may either be complementarity or substitutability between the bank's and the website's investment, depending on whether the bank's security measures are "active" or "passive".

Let us turn to the bank's problem. To focus the analysis is on the interactions between the bank and the website's security investments, I abstract away from the consumer's decision to use the bank by assuming that the consumer always remains with his bank, whether or not he detects a loss. This would correspond, for example, to a situation where the consumer faces high switching costs. This assumption allows us to reduce the bank's objective to the minimisation of the sum its investment cost and the fraud liability it would incur due to data breaches at the website[6]. The bank's loss function also depends on the nature of the security measures that it has in place. When its security measures are "passive", its loss function is given by

$$\phi_p(q^f,\gamma) = \begin{cases} (1-q^f)\lambda\rho_B((1+\delta)(1-\gamma)-\delta\lambda\rho_B(1-\gamma)^2)\alpha l & \text{if } \gamma \leq \overline{\gamma} \\ (1-q^f)(1+\delta)\lambda\rho_B(1-\gamma)\alpha l & \text{otherwise.} \end{cases}$$

When the measures are "active", its loss function is

$$\phi_a(q^f,\gamma) = \begin{cases} (1-q^f)\lambda\rho_B((1+\delta)(1-\gamma)-\delta\lambda\rho_B(1-\gamma)^2-\delta\rho_B\gamma(1-\gamma))\alpha l & \text{if } \gamma \leq \overline{\gamma} \\ (1-q^f)(1+\delta)\lambda\rho_B(1-\gamma)\alpha l & \text{otherwise.} \end{cases}$$

Let $t(\gamma)$ denote the amount the bank has to invest to attain a level of security, γ. Assume that the bank's investment cost function $t(\cdot)$ possesses the same properties as the website's. The bank's objective is to minimise its total cost

$$\psi_i(q^f,\gamma) = \phi_i(q^f,\gamma) + t(\gamma) \quad \text{where } i \in \{a,p\},$$

[6] For simplicity, I have assumed that the only source of losses for the bank arises from data breaches at the website. In reality, the bank may face other potential sources of losses (from the physical theft or loss of credit cards, for example) which may occur independently of data breaches at the website. The model can be easily modified to include these losses.

which gives us the following first-order conditions for "passive" and "active" measures respectively:

$$t'(\gamma_p^*) = \begin{cases} (1-q^f)\lambda\rho_B(1+\delta-2\delta\lambda\rho_B(1-\gamma_p^*))\alpha l & \text{if } \gamma \leq \overline{\gamma} \\ (1-q^f)(1+\delta)\lambda\rho_B\alpha l & \text{otherwise.} \end{cases} \tag{14}$$

and

$$t'(\gamma_a^*) = \begin{cases} (1-q^f)\lambda\rho_B(1+\delta-2\delta\lambda\rho_B(1-\gamma_a^*)+\delta\rho_B(1-2\gamma_a^*))\alpha l & \text{if } \gamma \leq \overline{\gamma} \\ (1-q^f)(1+\delta)\lambda\rho_B\alpha l & \text{otherwise.} \end{cases} \tag{15}$$

Consider the case of "passive" measures. A higher level of investment in security (when $\gamma \leq \overline{\gamma}$) generates opposing effects. On the one hand, it reduces the incidence of fraud following a breach. On the other hand, it promotes the use of an insecure website in the second period, since the consumer updates his belief favourably whenever no loss is detected. Because the second effect countervails the first, the optimal level of investment is lower than that associated with the simple minimisation of fraud losses. Indeed, in the case where the learning effect is irrelevant (i.e. when $\gamma > \overline{\gamma}$), one can observe that the marginal benefit of investment (and, hence, the level of investment) is always higher. By contrast, when the security measures are "active", the bank's investment promotes learning and lowers the likelihood an insecure website is used in the second period. The impact of the bank's investment on loss minimisation is reinforced by more consumer learning. Thus, the website has stronger incentive to invest with "active" rather than "passive" measures.

It is also worth noting that for both "active" and "passive" security measures, the bank's optimal level of security investment exhibits complementarity with the website's - the higher the level of investment made by the website, the lower its expected fraud liability and the lower its optimal level of investment.

5.1.2 Equilibrium Analysis

The equilibrium of this game can be found via backwards induction. Consider first the consumer's equilibrium strategy. Given the assumption made on his valuation as mentioned earlier, the consumer has a dominant strategy. When the level of protection offered by his bank is sufficiently high ($\gamma > \overline{\gamma}$), the consumer always uses the website. When $\gamma \leq \overline{\gamma}$, the consumer uses the website in both periods whenever he does not detect a loss in the first period and leaves the website in the second period otherwise.

Taking into account the consumer's strategy, the website and the bank simultaneously decide on their levels of investments. A Nash equilibrium of the investment game is given by the intersection between the website's and the bank's best response functions. From the website's first-order condition derived earlier, its best response function for "passive" bank security measures is given by:

$$q_p(\gamma_p) = \begin{cases} c'^{-1}(\delta\rho_B\lambda(1-\gamma_p)r) & \text{if } \gamma_p \leq \overline{\gamma} \\ 0 & \text{otherwise.} \end{cases}$$

The website's best response function is concave and downward-sloping for all values of $\gamma_p \in [0, \overline{\gamma})$. The more the bank invests, the less likely a fraud results from the breach and, hence, the less likely the consumer learns about the breach. This weakens the reputation effect and reduces the website's incentive to invest in security. It takes on a value of 0 for $\gamma_p \in [\overline{\gamma}, 1]$. The website has no incentive to invest since the consumer always uses the site regardless of its security level in this case.

The website's best response corresponding to the case of "active" bank security measures is

$$q_a(\gamma_a) = \begin{cases} c'^{-1}(\delta\rho_B(\lambda(1-\gamma_a)+\gamma_a)r) & \text{if } \gamma_a \leq \overline{\gamma} \\ 0 & \text{otherwise.} \end{cases}$$

In contrast to the case of "passive" measures, the website's best response function is concave and upward-sloping for all $\gamma_a \in [0, \overline{\gamma})$ for "active" measures. A higher the level of investment by the bank weakens the reputation effect when security measures are "passive" but it fosters learning (and strengthens the reputation effect) when the measures are "active". Thus, the website's incentive to invest is increasing with the bank's investment.

The derivation of the bank's best response function is slightly more complex. The bank's loss function is discontinuous in its level of investment, γ; this may result in multiple solutions to equations (14) and (15). In order to facilitate the analysis, I introduce the following lemma:

Lemma 2. Suppose that the bank's investment cost function t is such that

$$t'(\overline{\gamma}) > (1+\delta)\lambda\rho_B\alpha l.$$

Then, there exists only one solution $\gamma_i^*(q)$ (where $i \in \{a, p\}$) to the equations (15) and (14) respectively. Further, $\gamma_i^*(q) < \overline{\gamma}$ for all $q \in [0, 1]$.

The condition in Lemma 2 holds when the bank's investment cost is increasing rapidly with the level of investment; this corresponds to a high marginal cost of investment. Given the complex nature and the volume of financial transactions that a bank conducts, this may likely be the case. Whenever the condition in Lemma 2 is satisfied, the bank's loss function, and correspondingly, its best response function, are continuous in q. The bank's best response function is given by

$$\gamma_i(q) = \underset{\gamma}{\operatorname{argmin}}\ \psi(q, \gamma) \quad \text{where } i \in \{a, p\}.$$

It can be verified that the bank's best response function is a smooth, decreasing, and convex function of q for both "active" and "passive" measures.

Proposition 6. (Equilibrium in the Investment Game).
Suppose that t satisfies the condition in Lemma 2. There exists a unique, stable Nash

equilibrium in the security investment game. For "passive" security measures, the equilibrium levels of investment $\{q_p^*, \gamma_p^*\}$ satisfy:

$$\begin{cases} q_p^* = c^{-1}(\delta\rho_B\lambda(1-\gamma_p^*)r) \\ t'(\gamma_p^*) = (1-q_p^*)\lambda\rho_B(1+\delta-2\delta\lambda\rho_B(1-\gamma_p^*))\alpha l. \end{cases}$$

For "active" security measures, the equilibrium levels of investment $\{q_a^*, \gamma_a^*\}$ satisfy:

$$\begin{cases} q_a^* = c^{-1}(\delta\rho_B(\lambda(1-\gamma_a^*)+\gamma_a^*)r) \\ t'(\gamma_a^*) = (1-q_a^*)\lambda\rho_B(1+\delta-2\delta\lambda\rho_B(1-\gamma_a^*)+\delta\rho_B(1-2\gamma_a^*))\alpha l. \end{cases}$$

If the cost function t violates the condition in 2, the bank's best response function may have a kink at some value of q. Let us denote this level of investment $\underline{q}_i$ and note that, for the same set of parameter values, it differs for the two types of security measures. For all $q > \underline{q}_i$, the bank's optimal level of investment, $\gamma_i^*(q)$, lies below $\overline{\gamma}$; the converse is true for $q < \underline{q}_i$. At $q = \underline{q}_i$, the bank is indifferent between two levels of investment, $\gamma_i^*(\underline{q}_i)$ and $\gamma_i^{**}(\underline{q}_i)$ which lie above and below $\overline{\gamma}$ respectively. When this is the case, there could be a number of equilibrium outcomes, including one with mixed strategies. In order to facilitate the policy discussion, I will restrict the analysis to the class of bank investment cost functions for which Lemma 2 holds in the unregulated equilibrium (with $\lambda = \tilde{\lambda}$). Figures 9 to 12 provide graphical representations of two potential equilibrium outcomes discussed above for both "active" and "passive" bank security measures. Figures 10 and 12 illustrate how multiple equilibria can arise when the condition in lemma 2 is violated.

5.2 Policy Analysis

One important feature of the current framework is the interaction between the website's and the bank's investment decisions. The website's security investments acts as a strategic substitute to the bank's while the bank's investment, depending on the type of measures, may act as a complement or substitute to the website's investment. The failure to take into account these strategic interactions may lead to policy recommendations that lower the overall level of security. In this subsection, I will examine the regulatory implications of two policies - a unilateral bank security standard and mandatory data breach notification.

5.2.1 Unilateral Bank Security Standard

Although the level of fraud risk due to data breaches is determined by both the bank and the website, they do not belong to the same industry and may not be subject to the same set of regulations. Suppose that a regulator, seeking to improve the security of the payment system and reduce the incidence of fraud, imposes a unilateral minimum security standard, γ_{min}, on the bank. For example, the regulator could mandate the

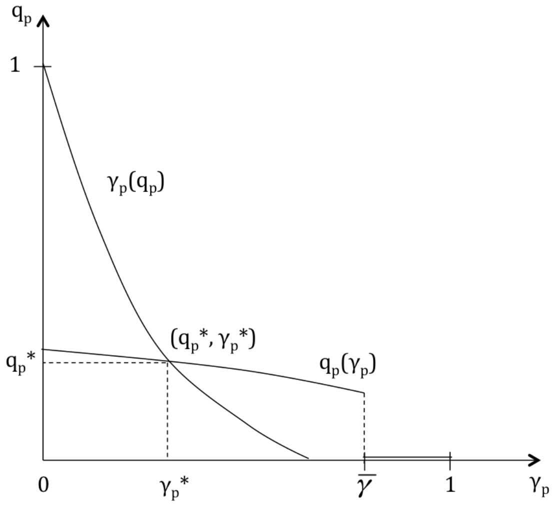

Figure 9: Unique equilibrium with “passive” bank security measures.

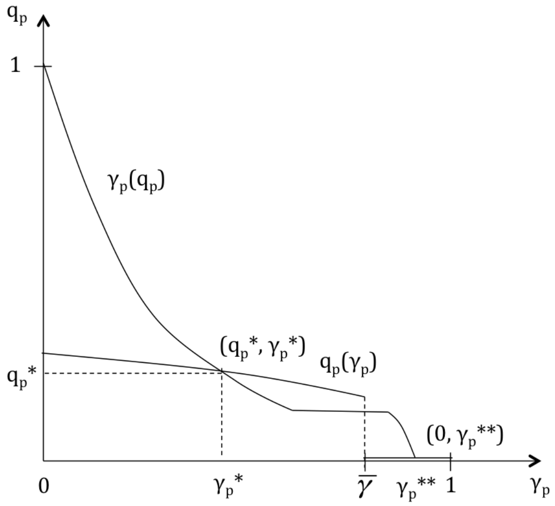

Figure 10: Two equilibria with “passive” bank security measures.

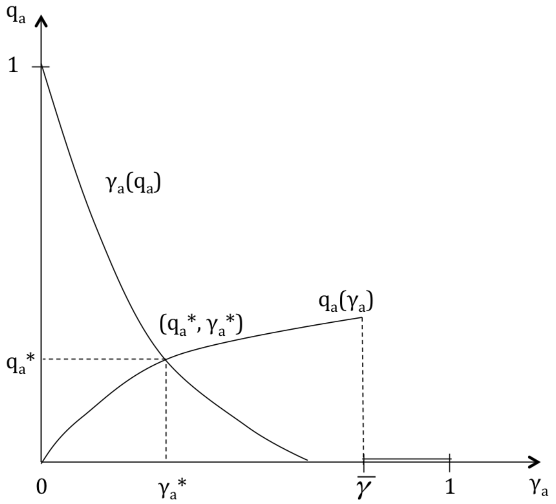

Figure 11: Unique equilibrium with “active” bank security measures.

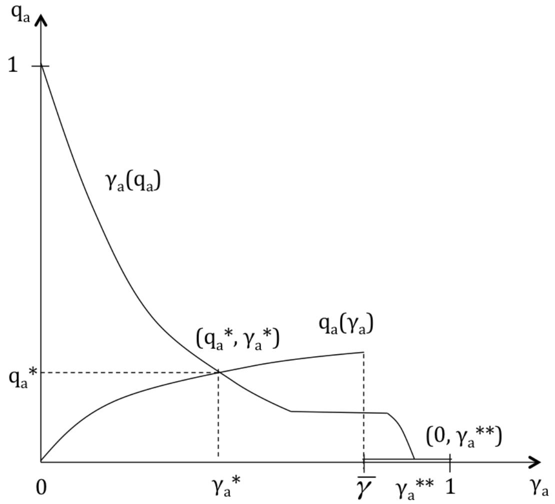

Figure 12: Two equilibria with “active” bank security measures.

bank to adopt ‘chip-and-PIN” payment cards. Would such a policy necessarily make the consumer better off?

Let us consider the case where the standard is higher than the unregulated equilibrium level of investment chosen by the bank (i.e. $\gamma_{min} > \gamma_i^*$), so that the standard has a bite. The impact of the minimum standard on consumer surplus depends on its effect on the overall level of security. Consumer surplus is always increasing in the overall level of security, since a higher level of security implies a lower probability of incurring fraud losses. Holding the website's level of investment constant, the standard would indeed raise the overall level of security (and, hence, consumer surplus). The policy's implications are, however, more ambiguous when we take into consideration the strategic behaviour of the website. This is summed up in the following proposition.

Proposition 7. (Minimum Security Standard on the Bank).
The effect of the minimum bank security standard on consumer surplus is
(i) positive when the bank's security measures are "active";
(ii) negative when the bank's security measures are "passive" if the following condition holds:

$$c''(q_p^*) < \delta \lambda \rho_B r \frac{1-\gamma_{min}}{1-q_{min}(\gamma_{min})}$$

(iii) ambiguous otherwise.

For "active" security measures, the consumer is notified whenever a fraud attempt is prevented by the bank. An increase in the bank's security level hence translates to an increase in the level of consumer learning. This strengthens the reputation effect and, thus, the website's incentive to invest. Since the standard raises both the bank's and the website's security levels, the consumer is unambiguously better off.

For "passive" security measures, the bank's investment in security acts as a substitute for the website's. The increase in the level of investment by the bank reduces the likelihood of a loss; consequently, the consumer is less likely to learn about the website's insecurity. This lowers the reputation cost associated with a breach, reducing the website's incentive to invest. Because the increase in bank's security is accompanied by a reduction in the website's security, the overall level of security could either increase or decrease. Whenever the condition in part (ii) of the proposition is satisfied, the decline in the website's security level dominates, lowering the overall level of security. Several factors affect the likelihood that the condition holds; one of which is the shape of the marginal cost function. Notice that the reduction in the website's investment is higher when the slope of the marginal cost function at q^* is less negative. When this is the case, a larger decrease in the level investment is needed to match the same level of decline in marginal benefit of investment (which results from the weaker reputation effect). The condition in the proposition is also more likely to hold when the mandated level of bank security, γ_{min}, is lower.

The proposition highlights two important points: first, the policy outcome depends on the type of security measures in question; second, the policy may be self-defeating when the website behaves strategically. Consider again the example of "chip-and-PIN" cards. Since "chip-and-PIN" cards are more difficult to counterfeit, the policy maker may believe that a regulation mandating the adoption of these cards will reduce the incidence of fraud and make the consumer better off. We know from the proposition, however, that this may not be true. "Chip-and-PIN" cards are a "passive" security measure. Although they reduce the probability of fraud, they also lower the likelihood of consumer learning. This weakens the website's incentive to invest and the overall effect on security may be ambiguous.

In view of the strategic interactions, a joint regulation of the bank and the website may be necessary in order to achieve the desired policy outcome.

5.2.2 Mandatory Breach Disclosure

Another policy of interest is the mandatory breach notification law that we discussed earlier. Much of the debate surrounding the proposed regulation has focused on its implications on firms' incentive to invest in security. Indeed, as we have seen that in the baseline model, mandatory breach notification increases the reputation cost of data breaches, which boosts the website's investment incentive. The potential impact of such a regulation on banks has, however, been overlooked. By raising the probability that the consumer learns about data breaches, a mandatory breach disclosure policy also increases the likelihood that the consumer detects fraudulent charges on his payment card. This generates a higher expected level of liability for the bank, which may in turn strengthen the bank's incentive to invest in security. The policy's impact on the bank's investment incentive depends on various factors, such as the initial level of fraud detection by the consumer and so on. These factors will also play a role in determining the investment levels of the website and bank in equilibrium with mandatory breach disclosure. In fact, the interplay between the website's and bank's incentive to invest and the underlying factors give rise many possible equilibrium outcomes. In order to have a more meaningful discussion, I will focus the analysis selectively on a few of the potential outcomes.

Suppose the breach disclosure regulation requires the website to notify the consumer of any breaches that have occurred, regardless of whether or not they have led to losses. Under this regulation, the website's profit function is the same for both "active" and "passive" security measures of the bank. Let π_{md} denote the website's profit under mandatory disclosure.

$$\pi_{md}(q^f, \gamma) = \begin{cases} (1+\delta-\delta(1-q^f)\rho_B)r - c(q^f) & \text{if } \gamma \leq \overline{\gamma}_{md} \\ (1+\delta)r - c(q^f) & \text{otherwise.} \end{cases}$$

This gives us the following best-response function:

$$q_{md}(\gamma_{md}) = \begin{cases} c'^{-1}(\delta\rho_B r) & \text{if } \gamma_{md} \leq \overline{\gamma}_{md} \\ 0 & \text{otherwise.} \end{cases}$$

$\overline{\gamma}_{md}$ corresponds to Equation (13) with $\lambda = 1$, that is

$$\overline{\gamma}_{md} = 1 - \frac{v}{\rho_B(1-\alpha)l}.$$

The regulation modifies both the interaction between the website and the bank and the website's incentive to invest. Notice that under the mandatory breach disclosure that

I have described, the website's profit and best-response function is independent of the bank's investment level, insofar as the bank's investment does not affect the consumer's participation decision. In the absence of the regulation, the consumer only learns if he detects a loss and/or is notified by the bank (for "active" measures) - both of which depends on the bank's security level. In the first scenario, the bank's security level affects the likelihood that a breach results in fraud and, consequently, the probability that the consumer learns about the breach. In the second scenario, the bank's security determines the frequency with which fraud attempts are being blocked and, hence, the probability that the consumer is notified. By mandating the disclosure of data breaches whether or not losses entail, the regulation renders both of these channels of learning redundant. One may also verify that the website's marginal benefit of investment is higher under the regulation relative to the unregulated case. The underlying intuition is the same as in the earlier set-up - mandatory disclosure makes a data breach more costly to the website by raising the probability of customer churn.

Mandatory disclosure also affects the amount of losses faced by the bank. When the consumer receives a notice that his data has been compromised, he is likely to check his bank statements for fraudulent transactions. This raises the probability that bank will be made liable for these transactions. For simplicity, suppose that mandatory disclosure translates to perfect detection of fraud. The bank's loss function is now the same regardless or whether its security measures are "active" or "passive" and is given by:

$$\phi_{md}(q^f, \gamma) = \begin{cases} (1-q^f)\rho_B(1+\delta-\delta\rho_B)(1-\gamma)\alpha l & \text{if } \gamma \leq \overline{\gamma}_{md} \\ (1-q^f)\rho_B(1+\delta)(1-\gamma)\alpha l & \text{otherwise.} \end{cases}$$

The corresponding first-order condition and best-response function are

$$t'(\gamma) = \begin{cases} (1-q^f)\rho_B(1+\delta-\delta\rho_B)\alpha l & \text{if } \gamma \leq \overline{\gamma}_{md} \\ (1-q^f)\rho_B(1+\delta)\alpha l & \text{otherwise.} \end{cases}$$

and

$$\gamma_{md}(q_{md}) = \begin{cases} t'^{-1}((1-q_{md})\rho_B(1+\delta-\delta\rho_B)\alpha l) & \text{if } \gamma_{md} \leq \overline{\gamma}_{md} \\ t'^{-1}((1-q_{md})\rho_B(1+\delta)\alpha l) & \text{otherwise.} \end{cases}$$

Mandatory breach notification generates two opposing effects - fraud detection and consumer learning - on the bank's investment incentive to invest. First, by increasing the likelihood of fraud detection by the consumer, mandatory disclosure raises the expected liability that the bank has to assume. This creates a stronger incentive for the bank to invest. Second, breach notification leads to consumer learning, which decreases that likelihood that the consumer uses an insecure website in the second period. This lowers the expected fraud liability cost for the bank, hence reducing the benefit of investment. The relative strengths of these two effects depend on the initial level of loss detection, $\tilde{\lambda}$. For low initial levels of detection ($\tilde{\lambda}$ is small), the former effect dominates and the bank's incentive to invest is bolstered by the regulation. The following proposition summarises the above discussion.

Proposition 8. (Increase in Investment by both Bank and Website).
Consider the case where $\tilde{\lambda}$ is sufficiently small. Mandatory breach notification raises both the website's and the bank's security investment in equilibrium. The equilibrium levels of investment satisfy:

$$\begin{cases} q^*_{md} = c'^{-1}(\delta\rho_B r) \\ \gamma^*_{md} = t'^{-1}((1 - q^*_{md})\rho_B(1 + \delta - \delta\rho_B)\alpha l). \end{cases}$$

When $\tilde{\lambda}$ is small, the regulation raises both the bank's and the website's levels of investment, which translates into a higher overall level security. The consumer is less likely to incur fraud losses and is unambiguously better off. The website is made worse off - customer turnover is more likely despite the higher level of security investment. The effect on the bank's overall cost is ambiguous.

The impact of mandatory disclosure on the equilibrium level of investment may be less certain in the case where $\tilde{\lambda}$ is large. The regulation lowers the bank's optimal level of investment for a range of the website's investment level, q. As a result, the equilibrium could feature a higher level of investment made by the website but a lower one made by the bank. The overall impact of the policy would depend on the relative magnitudes of the change in investment levels - consumer may be made better or worse off (depending on the resulting overall security level). As before, the website is made worse off and the bank may be either better or worse off.

Another interesting policy outcome is characterised by severe consumer moral hazard and the full crowding-out of the website's investment. In the equilibrium outcomes we have considered so far, the bank's optimal investment level lies below $\overline{\gamma}_{md}$. It is possible, however, that the cost minimising level of investment γ_{md} exceeds $\overline{\gamma}_{md}$ under the regulation. Should this be the case, the consumer would always participate regardless of the website's security level. Consequently, the website would have no incentive to invest in security. The following proposition formalises this result.

Proposition 9. (Full Crowding Out of Website's Investment).
Suppose the the cost function t satisfies the following condition:

$$t'(\overline{\gamma}_{md}) < (1 - c'^{-1}(\delta\rho_B r))\rho_B(1 + \delta - \delta\rho_B)\alpha l.$$

Then, the equilibrium levels of investment are given by

$$\begin{cases} q^*_{md} = 0 \\ \gamma^*_{md} = t'^{-1}((1 + \delta)\rho_B\alpha l). \end{cases} \tag{16}$$

The website's investment is completely crowded out by the increase in the bank's. The high level of investment made by the bank caps the expected losses that the consumer faces when using an insecure website and results in indiscriminate usage. This equilibrium outcome is likely when the website' incentive to invest is generally low. This could correspond to the case where the website's revenue, r, is relatively small, and hence, the reputation cost of a breach is small. It could also be that the website cares little about its future revenue flow (for example, if the website believes that it is not likely to continue operating in the next period); that is, δ for the website is small [7].

When the equilibrium outcome is as described in Equation (16), both the consumer and the website are made better off, while the bank is made worse off. This is clearly not optimal from the society's point of view, however. The overall cost to the society would be lower if some security investment was undertaken by the website, whose marginal cost of investment is zero at $q = 0$. One means of shifting some of the investment burden to the website is to make it liable for part of the breach induced losses incurred by the bank. This could potentially be achieved by implementing a well-defined liability rule or loss-shifting arrangement. A negligence rule may also be introduced to address the consumer moral hazard problem. The design of an optimal joint regulatory framework is left to future work.

6 Conclusion

In this paper, I proposed a model of data security investment in which the firm's incentive to invest is driven by the reputation effect of data breaches. It represents the first attempt to model data security as an unobserved endogenous quality dimension of a firm. The baseline model comprises of a website and a unit mass of consumers. The website makes an unobserved one-time security investment and the consumers learn about the firm's security over time via imperfect breach detection. The extended model further includes the consumer's bank. There, the overall level of security is jointly determined by the website and the bank. This generates interdependencies in security investments.

Underinvestment in security arises in this model because the firm does not fully internalise the losses that data breaches impose on the society (the consumers and their banks). Further, information asymmetries between the firm and its consumers create a moral hazard problem. I examine two policies - mandatory breach notification and a minimum bank security standard - that may be introduced to improve the overall level of security. Mandatory breach notification aims at strengthening the firm's investment incentive by increasing the reputation cost of data breaches, while a minimum bank security standard raises the bank's security level directly. I demonstrate, however, that these policies need not necessarily lead to the desired improvement in security (and may even be self-defeating), when we take into account the strategic behaviours of the agents.

I have made several simplifying assumptions in this paper. First, the website makes a one-time investment in security in this model. It may also be interesting to allow for the website to re-invest in security, particularly after a breach. Second, I have

[7] I have assumed that the discount factors of the bank and the website are the same in the analysis but the results would remain unchanged even when we allow for individual-specific discount factors.

assumed the security threat that the firm faces to be exogenous. As an extension, one could also examine the case where the probability of attack depends on the firm's level of investment (for example, see Hausken (2006) and Cavusoglu et al. (2008)) or that where it is contingent on the size of the firm's customer base. Third, I focused on the case where there is a single firm. When consumers use multiple firms, it may become difficult for them to ascertain the source of a data breach, even if they were to discover a fraudulent transaction. Under this scenario, the consumers' bank may play a bigger role as a monitor of firms' security. By observing pattern of fraud across the consumers' accounts, the bank may be able to identify the firm linked to the data breach[8]. The monitoring role of the bank may be studied in greater detail in future work. Finally, I have considered the case where security investments (by both the website and the bank) are made at the beginning of the game. In doing so, I have excluded the possibility that the bank may make a "supplementary investment" in security in response to breach announcements. For example, the bank may reissue payment cards that have been exposed but have not (yet) been compromised. The optimality of card reissuance have been examined by Graves et al. (2014), who compared the cost of reissuing with that of not reissuing (i.e. potential fraud losses) via Monte Carlo simulations. In addition to the direct costs, the optimality of re-issuance may also depend on its incentive effects on the firm. Future work could explore these incentive effects by extending the model to allow for the bank to make a "supplementary investment" in response to a breach.

[8] A recent case of credit card breach at the fastfood chain, Wendy's, provides an example of the role of the banking industry as a monitor. See article: `http://krebsonsecurity.com/2016/01/wendys-probes-reports-of-credit-card-breach/`.

References

Acquisti, A. and J. Grossklags (2005): "Privacy and rationality in individual decision making," *IEEE Security & Privacy*, 2, 24–30.

American Bankers Association (2016, January 2016): "Banks Stop $11 Billion in Fraud Attempts in 2014," http://www.aba.com/Press/Pages/012716DepositSurvey.aspx.

Anderson, R. and T. Moore (2006): "The economics of information security," *Science*, 314, 610–613.

Board, S. and M. Meyer-ter Vehn (2013): "Reputation for quality," *Econometrica*, 81, 2381–2462.

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

Cavusoglu, H., B. Mishra, and S. Raghunathan (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, 70–104.

Cavusoglu, H., S. Raghunathan, and W. T. Yue (2008): "Decision-theoretic and game-theoretic approaches to IT security investment," *Journal of Management Information Systems*, 25, 281–304.

Gordon, L. A. and M. P. Loeb (2002): "The economics of information security investment," *ACM Transactions on Information and System Security (TISSEC)*, 5, 438–457.

Gorman, G. and P. Evan (2009): "Hackers indicted in widespread ATM heist," *The Wall Street Journal*, November 11, Available: http://www.wsj.com/articles/SB125786711092441245.

Graves, J. T., A. Acquisti, and N. Christin (2014): "Should payment card issuers reissue cards in response to a data breach," in *Proceedings of the 2014 Workshop on the Economics of Information Security.*

Grossklags, J., N. Christin, and J. Chuang (2008): "Secure or insecure? A game-theoretical analysis of information security games," in *Proceedings of the 17th International World Wide Web Conference*, 209–218.

Hausken, K. (2006): "Income, interdependence, and substitution effects affecting incentives for security investment," *Journal of Accounting and Public Policy*, 25, 629–665.

Kerber, R. (2013): "Target payment card data theft highlights lagging U.S. security," *Reuters*, December 22, Available: `http://www.reuters.com/article/us-target-security-lagging-idUSBRE9BL06X20131222#HQ6Uckr8BkZvsAQo.97`.

Kunreuther, H. and G. Heal (2003): "Interdependent security," *Journal of risk and uncertainty*, 26, 231–249.

Risk Based Security (2015): "Data breach QuickView: 2015 data breach trends," .

Romanosky, S., R. Sharp, and A. Acquisti (2010): "Data Breaches and Identity Theft: When is Mandatory Disclosure Optimal?" .

Romanosky, S., R. Telang, and A. Acquisti (2011): "Do data breach disclosure laws reduce identity theft?" *Journal of Policy Analysis and Management*, 30, 256–286.

Shapiro, C. (1983): "Premiums for high quality products as returns to reputations," *The quarterly journal of economics*, 659–679.

Sullivan, R. J. (2010): "The changing nature of US card payment fraud: industry and public policy options," *Economic Review-Federal Reserve Bank of Kansas City*, 95, 101.

The New York Times (2010, August 21): "$9 Here, 20 Cents There and a Credit-Card Lawsuit," `http://www.nytimes.com/2010/08/22/business/22digi.html?_r=1`.

Varian, H. (2004): "System reliability and free riding," in *Economics of information security*, Springer, 1–15.

Appendix

Bayesian Consumers with Exogenous Beliefs

In the paper, I have assumed that consumers' have rational beliefs which are consistent with the equilibrium level of investment chosen by the firm (i.e. $q^c = q^*$). Here, I will consider the case where the consumers have inconsistent exogenous beliefs.

Let $\tilde{q}$ denote the optimal level of investment when consumers have exogenous beliefs. For any given level of consumer belief, $\tilde{q}$ is pinned down by

$$c'(\tilde{q}) = \delta\lambda\rho_B \left(\frac{(1+\delta)q^c}{1+\delta(1-\lambda(1-q^c)\rho_B)} \right) \rho_B Lr.$$

The right-hand side of the expression gives us the marginal benefit of security investment. One can easily verify that this is increasing in the initial belief, q^c, of the consumers. In other words, the better the initial reputation of the website, the stronger its incentive to invest. The intuition underlying this result is simple. A website with a better reputation attracts more consumers in the first period, which also implies that it has more business to lose when a breach occurs.

We can distinguish between two possible scenarios - consumers may either be over-optimistic ($q^c > q^*$) or over-pessimistic ($q^c < q^*$). In the over-optimism scenario, the website's investment is higher than in the market equilibrium with rational beliefs; more specifically, we have that $q^* < \tilde{q} < q^c$. In the over-pessimism scenario, we obtain $q^c < \tilde{q} < q^*$. Consumers' beliefs are, to some extent, self-fulfilling: higher beliefs lead to higher levels of investment and vice-versa.

Relative to the case of rational consumers, the website is better off when consumers are over-optimistic (and worse off otherwise). The effect of over-optimism on consumer surplus is more ambiguous. On the one hand, it raises the level of security investment, which reduces the expected loss from using the website. On the other hand, it leads to excessive usage of the website. That said, over-optimism is more likely to benefit the consumers when the marginal cost of investment raises less rapidly. For any initial level of belief, the website's optimal level of investment is higher when that is the case.

Proof of Proposition 1

The equilibrium level of q in this model satisfies Equation (11):

$$c'(q) = \delta\lambda\rho_B^2 \left(\frac{(1+\delta)q}{1+\delta(1-\lambda(1-q)\rho_B)} \right) Lr.$$

Let us refer to the right hand side of this expression as the equilibrium marginal benefit function and denote it $MB(\cdot)$.

I will first show that $MB(\cdot)$ is increasing and concave for $r > 0$. For all $q \in [0,1]$,

$$\frac{\partial MB(q)}{\partial q} = \delta\lambda\rho_B^2 \left(\frac{(1+\delta)(1+\delta(1-\lambda\rho_B))}{(1+\delta(1-\lambda(1-q)\rho_B))^2} \right) Lr > 0,$$

which indicates that $MB(\cdot)$ is increasing and

$$\frac{\partial^2 MB(q)}{\partial q^2} = -\delta^2\lambda^2\rho_B^3 \left(\frac{2(1+\delta)(1-\delta(1-\lambda\rho_B))}{(1+\delta(1-\lambda(1-q)\rho_B))^3} \right) Lr < 0,$$

which means that $MB(\cdot)$ is concave.

Let us define the $H(\cdot)$ as a function of the difference between $MB(\cdot)$ and $c'(\cdot)$:

$$H(q) = MB(q) - c'(q).$$

Since both $MB(\cdot)$ and $c'(\cdot)$ are continuous in q for $q \in [0,1]$, $H(.)$ is also continuous in q on the same interval.

Consider the case where $c''(0) = 0$. We have that

$$H(0) = 0 - 0 = 0 \quad \text{and} \quad H(1) = \delta\lambda\rho_B^2 Lr - \infty = -\infty.$$

Further, we know that

$$H'(0) = \delta\lambda\rho_B^2 \left(\frac{1+\delta}{1+\delta(1-\lambda\rho_B)} \right) Lr - 0 > 0 \quad \text{if } r > 0$$

This implies that there exists a $q_0 \in (0,1]$ such that $H(q_0) > 0$.

Since $H(\cdot)$ is continuous on the interval $[q_0, 1]$, by the Intermediate Value Theorem (IVT), there exist a value $q^* \in [q_0, 1]$ such that $H(q^*) = 0$. Thus, a positive investment equilibrium always exists.

Proof of Proposition 2

(i) For any arbitrary level of consumer beliefs and website's choice of investment level, consumer surplus can be decomposed as follows:

$$CS(q^f, q^c) = q^f CS^(1, q^c) + (1-q^f)CS(0, q^c).$$

The first order derivative of consumer surplus with respect to the website's investment policy is

$$\frac{\partial CS(q^f, q^c)}{\partial q^f} = CS(1, q^c) - CS(0, q^c) > 0.$$

Thus, for any given level of consumer beliefs, consumers are always better if the website chooses a higher level of q^f. In particular, this implies that

$$CS(0,0) < CS(q_+^*, 0).$$

Moreover, by revealed preference

$$CS(q_+^*, 0) < CS(q_+^*, q_+^*).$$

Therefore, we have that

$$CS(0,0) < CS(q_+^*, 0) < CS(q_+^*, q_+^*),$$

that is, the consumer surplus at the positive investment equilibrium is higher than that under the no investment equilibrium.

I will now show that $\pi(0,0) \leq \pi(q_+^*, q_+^*)$. First, by revealed preference, we have that

$$\pi(q, q_+^*) \leq \pi(q_+^*, q_+^*) \quad \forall q \neq q_+^*.$$

In particular,

$$\pi(0, q_+^*) \leq \pi(q_+^*, q_+^*).$$

In addition,

$$\pi(0, q_+^*) - \pi(0,0) = (1+\delta)\left(\hat{v}(0) - \hat{v}(q_+^*)\right) r + \delta\lambda\rho_B\left(\hat{v}(q_+^*) - \hat{v}_D\right).$$

Since $\hat{v}(0) = \hat{v}_D$, the above expression can be rewritten as

$$\pi(0, q_+^*) - \pi(0,0) = (1+\delta(1-\lambda\rho_B))\left(\hat{v}_D - \hat{v}(q_+^*)\right) r > 0.$$

Thus, we have that

$$\pi(0,0) < \pi(0, q^*) \leq \pi(q^*, q^*).$$

Since $CS(0,0) < CS(q^*, q^*)$ and $\pi(0,0) < \pi(q^*, q^*)$, the positive investment equilibrium Pareto dominates the no investment equilibrium.

(ii) I will now show that the website always under-invests relative to the social optimum. Recall that social welfare function is comprised of two components: the website's profit and consumer surplus.

$$SW(q,q) = \pi(q,q) + CS(q,q).$$

The change in social welfare corresponding to an increase in q at the private equilibrium is given by

$$\left.\frac{\partial SW(q,q)}{\partial q}\right|_{q=q^*} = \left.\frac{\partial \pi(q,q)}{\partial q}\right|_{q=q^*} + \left.\frac{\partial CS(q,q)}{\partial q}\right|_{q=q^*},$$

with $q^* = \{0\}$ when only the no investment equilibrium exists and $q^* = \{0, q_+^*\}$ when both equilibria exist.

I will first show that the website's profit in equilibrium is increasing in q^*.

$$\left.\frac{\partial \pi(q,q)}{\partial q}\right|_{q=q^*} = -1 + \delta(1-\lambda(1-q^*)\rho_B)\left(\left.\frac{\partial \hat{v}(q)}{\partial q}\right|_{q=q_+^*}\right) r,$$

where

$$\left.\frac{\partial \hat{v}(q)}{\partial q}\right|_{q=q_+^*} = -\left(\frac{(1+\delta)(1+\delta(1-\lambda\rho_B))}{(1+\delta(1-(1-q^*)\lambda\rho_B))^2}\right)\rho_B L < 0.$$

Hence,

$$\left.\frac{\partial \pi(q,q)}{\partial q}\right|_{q=q^*} > 0.$$

Further, I have demonstrated in the proof for part (i) that consumer surplus is increasing in the website's level of investment. Hence,

$$CS(q', q) \geq CS(q, q) \quad \forall q' \geq q.$$

By the optimality of consumers' decision, we also know that

$$CS(q', q') \geq CS(q', q).$$

This gives

$$CS(q', q') \geq CS(q', q) \geq CS(q, q).$$

Thus,

$$\left.\frac{\partial SW(q,q)}{\partial q}\right|_{q=q^*} \geq 0.$$

and the social planner would also prefer a higher level of investment than the website. In other words, the website under-invests relative to the social optimum.

Proof of Proposition 3

Let us denote the marginal benefit of investment under the individual notice requirement as $MB_{in}(q)$. It corresponds to the marginal benefit function with $\lambda = 1$:

$$MB_{in}(q) = \delta\rho_B\left(\frac{(1+\delta)q}{1+\delta(1-(1-q)\rho_B)}\right)\rho_B(1-\alpha)lr$$

Mandatory breach notification induces a higher level of investment in equilibrium when the marginal benefit of investing in the presence of breach notification exceeds that of the case without the regulation. More specifically, if $MB_{in}(q^*) > MB(q^*)$, the firm always invests more when individual notice is required ($q^*_{in} > q^*$). At any level of q,

$$\begin{aligned} MB_{in}(q) - MB(q) = {} & \delta\rho_B\left(\frac{(1+\delta)q}{1+\delta(1-(1-q)\rho_B)}\right)\rho_B(1-\alpha)lr \\ & - \delta\tilde{\lambda}\rho_B\left(\frac{(1+\delta)q}{1+\delta(1-\tilde{\lambda}(1-q)\rho_B)}\right)\rho_B(1-\tilde{\lambda}\alpha)lr \end{aligned}$$

and

$$\begin{aligned} \frac{\partial MB_{in}(q)}{\partial q} - \frac{\partial MB(q)}{\partial q} = {} & \delta\rho_B\left(\frac{(1+\delta)(1+\delta(1-\rho_B))}{(1+\delta(1-(1-q)\rho_B))^2}\right)\rho_B(1-\alpha)lr \\ & - \delta\tilde{\lambda}\rho_B\left(\frac{(1+\delta)(1+\delta(1-\tilde{\lambda}\rho_B))}{(1+\delta(1-\tilde{\lambda}(1-q)\rho_B))^2}\right)\rho_B(1-\tilde{\lambda}\alpha)lr. \end{aligned}$$

For MB_{in} to lie everywhere above MB, we require that
(i) The slope of MB_{in} is larger than that of MB at $q = 0$, i.e.

$$\begin{aligned} & \left.\frac{\partial MB_{in}(q)}{\partial q}\right|_{q=0} - \left.\frac{\partial MB(q)}{\partial q}\right|_{q=0} > 0 \\ \Leftrightarrow \quad & \delta\rho_B\left(\frac{1+\delta}{1+\delta(1-\rho_B)}\right)\rho_B(1-\alpha)lr \\ & - \delta\tilde{\lambda}\rho_B\left(\frac{1+\delta}{1+\delta(1-\tilde{\lambda}\rho_B)}\right)\rho_B(1-\tilde{\lambda}\alpha)lr > 0 \\ \Leftrightarrow \quad & \alpha < \frac{1+\delta}{1+\delta+\tilde{\lambda}(1+\delta(1-\rho_B))}. \end{aligned}$$

(ii) The value of MB_{in} is larger than that of MB at $q = 1$.

$$\begin{aligned} & MB_{in}(1) - MB(1) > 0 \\ \Leftrightarrow \quad & \delta\rho_B^2(1-\alpha)lr - \delta\tilde{\lambda}\rho_B^2(1-\tilde{\lambda}\alpha)lr > 0 \\ \Leftrightarrow \quad & \alpha < \frac{1}{1+\tilde{\lambda}} \end{aligned}$$

It can be verified that $\frac{1}{1+\tilde{\lambda}} < \frac{1+\delta}{1+\delta+\tilde{\lambda}(1+\delta(1-\rho_B))}$ Thus, MB_{in} lies above MB for all $q > 0$ when

$$0 \le \alpha < \frac{1}{1+\tilde{\lambda}}.$$

By the same token, when conditions (i) and (ii) are violated, MB_{in} must lie everywhere below MB. This corresponds to

$$\alpha > \frac{(1+\delta)}{(1+\delta+\tilde{\lambda}(1+\delta(1-\rho_B))}.$$

For a given increasing marginal cost function c', it is clear that if MB_{in} lies everywhere above MB, MB_{in} must intersect c' at a higher value of q than MB. Hence, we have that $q^*_{in} > q^*$. The converse is true when MB_{in} lies everywhere below MB.

For $\dfrac{1}{1+\tilde{\lambda}} \leq \alpha < \dfrac{(1+\delta)}{(1+\delta+\tilde{\lambda}(1+\delta(1-\rho_B))}$, the effect of increasing λ from $\tilde{\lambda}$ to 1 is less certain. MB_{in} lies above MB for some values of q and below for others. In fact, for any given set of parameter values, there exists a $\hat{q}$ such that for all $q < \hat{q}$, MB_{in} lies above MB. Define $J(\cdot) = MB_{in}(\cdot) - MB(\cdot)$; $J(\cdot)$ is continuous for $q \in [0,1]$. For α within the specified range, $J(q) > 0$ for small q and $J(q) < 0$ for large q. By the IVT, there exists a $\hat{q}$ such that $MB_{in}(\hat{q}) = MB(\hat{q})$ and we have that $MB_{in}(q) > MB(q)$ for all $q < \hat{q}$ and vice-versa. Further, since $c'(q)$ is increasing in q, for q^*_{in} to be higher than q^*, we need to have $MB_{in}(q^*) \geq MB(q^*) = c(q^*)$. This is equivalent to

$$r \geq c'(q^*)\left(\frac{(1+\delta(1-(1-q^*)\rho_B))}{\delta\rho_B^2(1+\delta)q^*(1-\alpha)l}\right).$$

Proof of Proposition 4

The proof of this proposition follows in the same line of reasoning as that of the previous proposition. Let us denote the marginal benefit of investment under the media notice requirement as $MB_{mn}(q)$.

$$MB_{mn}(q) = \delta\rho_B\left(\frac{q}{1-(1-q)\rho_B}\right)\rho_B(1-\alpha)lr.$$

We have that

$$\begin{aligned} MB_{mn}(q) - MB(q) &= \delta\rho_B\left(\frac{q}{1-(1-q)\rho_B}\right)\rho_B(1-\alpha)lr. \\ &\quad - \delta\tilde{\lambda}\rho_B\left(\frac{(1+\delta)q}{1+\delta(1-\tilde{\lambda}(1-q)\rho_B)}\right)\rho_B(1-\tilde{\lambda}\alpha)lr \end{aligned}$$

and

$$\begin{aligned} \frac{\partial MB_{mn}(q)}{\partial q} - \frac{\partial MB(q)}{\partial q} &= \delta\rho_B\left(\frac{(1-\rho_B)}{(1-(1-q)\rho_B)^2}\right)\rho_B(1-\alpha)lr \\ &\quad - \delta\tilde{\lambda}\rho_B\left(\frac{(1+\delta)(1+\delta(1-\tilde{\lambda}\rho_B))}{(1+\delta(1-\tilde{\lambda}(1-q)\rho_B))^2}\right)\rho_B(1-\tilde{\lambda}\alpha)lr. \end{aligned}$$

For MB_{mn} to lie everywhere above MB, we require that

(i) The slope of MB_{mn} is larger than that of MB at $q = 0$, i.e.

$$\begin{aligned}
& \left.\frac{\partial MB_{mn}(q)}{\partial q}\right|_{q=0} - \left.\frac{\partial MB(q)}{\partial q}\right|_{q=0} > 0 \\
\Leftrightarrow \quad & \delta\rho_B \left(\frac{1}{1-\rho_B}\right) \rho_B(1-\alpha)lr \\
& - \delta\tilde{\lambda}\rho_B \left(\frac{1+\delta}{1+\delta(1-\tilde{\lambda}\rho_B)}\right) \rho_B(1-\tilde{\lambda}\alpha)lr > 0 \\
\Leftrightarrow \quad & \alpha < \frac{(1+\delta)(1-\tilde{\lambda})+\tilde{\lambda}\rho_B}{(1+\delta)(1-\tilde{\lambda}^2)-\delta\tilde{\lambda}\rho_B(1-\tilde{\lambda})+\tilde{\lambda}^2\rho_B}.
\end{aligned}$$

(ii) The value of MB_{mn} is larger than that of MB at $q = 1$.

$$\begin{aligned}
& MB_{mn}(1) - MB(1) > 0 \\
\Leftrightarrow \quad & \delta\rho_B^2(1-\alpha)lr - \delta\tilde{\lambda}\rho_B^2(1-\tilde{\lambda}\alpha)lr > 0 \\
\Leftrightarrow \quad & \alpha < \frac{1}{1+\tilde{\lambda}}
\end{aligned}$$

It can be verified that $\frac{1}{1+\tilde{\lambda}} < \frac{(1+\delta)(1-\tilde{\lambda})+\tilde{\lambda}\rho_B}{(1+\delta)(1-\tilde{\lambda}^2)-\delta\tilde{\lambda}\rho_B(1-\tilde{\lambda})+\tilde{\lambda}^2\rho_B}$. Thus, MB_{mn} lies above MB for all $q > 0$ when

$$0 \leq \alpha < \frac{1}{1+\tilde{\lambda}}.$$

By the same token, when conditions (i) and (ii) are violated, MB_{mn} must lie everywhere below MB. This corresponds to

$$\alpha > \frac{(1+\delta)(1-\tilde{\lambda})+\tilde{\lambda}\rho_B}{(1+\delta)(1-\tilde{\lambda}^2)-\delta\tilde{\lambda}\rho_B(1-\tilde{\lambda})+\tilde{\lambda}^2\rho_B}.$$

For a given increasing marginal cost function c', if MB_{in} lies everywhere above MB, MB_{mn} must intersect c' at a higher value of q than MB. Hence, we have that $q^*_{mn} > q^*$. The converse is true when MB_{mn} lies everywhere below MB.

For $\frac{1}{1+\tilde{\lambda}} \leq \alpha < \frac{(1+\delta)(1-\tilde{\lambda})+\tilde{\lambda}\rho_B}{(1+\delta)(1-\tilde{\lambda}^2)-\delta\tilde{\lambda}\rho_B(1-\tilde{\lambda})+\tilde{\lambda}^2\rho_B}$, there exists a $\hat{q}$ such that for all $q < \hat{q}$, MB_{mn} lies above MB. Further, since $c'(q)$ is increasing in q, for q^*_{mn} to be higher than q^*, we need to have $MB_{mn}(q^*) \geq MB(q^*) = c(q^*)$. This is equivalent to

$$r \geq c'(q^*)\left(\frac{1-(1-q^*)\rho_B}{\delta\rho_B^2 q^*(1-\alpha)l}\right).$$

To show that $q^*_{mn} \geq q^*_{in}$, we need to establish that $MB_{mn}(q) > MB_{in}(q)$ for all q. Indeed, we have that

$$\begin{aligned} MB_{mn}(q) - MB_{in}(q) &= \delta \rho_B \left(\frac{q}{1-(1-q)\rho_B} \right) \rho_B (1-\alpha) lr \\ &- \delta \rho_B \left(\frac{(1+\delta)q}{1+\delta(1-(1-q)\rho_B))} \right) \rho_B (1-\alpha) lr \\ &\propto \frac{(1-q)\rho_B}{(1-(1-q)\rho_B)(1+\delta(1-(1-q)\rho_B)} \\ &\geq 0. \end{aligned}$$

Thus, we can conclude that $q^*_{mn} \geq q^*_{in}$.

Proof of Proposition 5

(i) Social welfare is obtained as the sum of consumer surplus and the website's profit. Let $CS^{in}(\cdot,\cdot)$ denote the consumer surplus function under mandatory individual notice; it corresponds to $CS(\cdot,\cdot)$ but with $\lambda = 1$. I will first show that consumer surplus is higher under breach notification whenever it leads to a higher level of investment ($q^*_{in} \geq q^*$), i.e.

$$CS^{in}(q^*_{in}, q^*_{in}) \geq CS(q^*, q^*).$$

We have established in the proof of proposition 3 that for any given level of consumer belief, consumer surplus is increasing in the website's level of investment. Hence,

$$CS(q^*_{in}, q^*) \geq CS(q^*, q^*).$$

Further, the optimality of the consumers' decision implies

$$CS(q^*_{in}, q^*_{in}) \geq CS(q^*_{in}, q^*).$$

In addition, one can show an increase in λ results in a higher level of consumer surplus ceteris paribus (i.e. $\dfrac{\partial CS(q,q)}{\partial \lambda} > 0$ for all $q \in [0,1]$). Therefore,

$$CS_{in}(q^*_{in}, q^*_{in}) \geq CS(q^*_{in}, q^*_{in}).$$

Putting together the above inequalities, we obtain

$$CS_{in}(q^*_{in}, q^*_{in}) \geq CS(q^*_{in}, q^*_{in}) \geq CS(q^*_{in}, q^*) \geq CS(q^*, q^*).$$

Thus, we have shown that $CS_{in}(q^*_{in}, q^*_{in}) \geq CS(q^*, q^*)$ whenever $q^*_{in} \geq q^*$.

I will now show that the website's profit is also higher under mandatory individual notification whenever $q^*_{in} \geq q^*$. Let $\pi_{in}(\cdot,\cdot)$ denote the profit function of the website under the regulation. $\pi_{in}(\cdot,\cdot)$ corresponds to $\pi(\cdot,\cdot)$ with $\lambda = 1$. By the same token, let $\hat{v}^n(\cdot)$ and $\hat{v}^n_D$ denote the usage thresholds when $\lambda = 1$.

By the optimality of the website's investment decision, we have that

$$\pi_{in}(q^*_{in}, q^*_{in}) \geq \pi_{in}(q^*, q^*_{in}).$$

Further, since $\hat{v}(q^*) \geq \hat{v}^n(q^*_{in})$ and $\hat{v}_D \geq \hat{v}^n_D$ whenever $q^*_{in} \geq q^*$, we have that

$$\begin{aligned}\pi_{in}(q^*, q^*_{in}) - \pi(q^*, q^*) &= (1 + \delta(1 - (1 - q^*))\rho_B)(\hat{v}(q^*) - \hat{v}^n(q^*_{in}))r \\ &\quad + \delta(1 - q^*)[\rho_B(\hat{v}_D - \hat{v}^n_D) + (1 - \lambda)(\hat{v}_D - \hat{v}(q^*))]r \\ &\geq 0.\end{aligned}$$

Hence,

$$\pi_{in}(q^*, q^*_{in}) \geq \pi(q^*, q^*).$$

Combining the two inequalities, we obtain

$$\pi_{in}(q^*_{in}, q^*_{in}) \geq \pi_{in}(q^*, q^*_{in}) \geq \pi(q^*, q^*).$$

Therefore, we have established that $\pi_{in}(q^*_{in}, q^*_{in}) \geq \pi(q^*, q^*)$ whenever $q^*_{in} \geq q^*$.

Since both consumer surplus and profit are higher under mandatory individual notification whenever $q^*_{in} \geq q^*$, social welfare is unambiguously higher.

(ii) I will adopt the same approach to establish that social welfare is always higher with the additional requirement of media notice. Let $CS_{mn}(\cdot,\cdot)$ and $\pi_{mn}(\cdot,\cdot)$ denote the consumer surplus and profit functions with mandatory media notice and let $\hat{v}^n_M$ and $\hat{v}^n_{ND}$ denote the usage thresholds in the first period for myopic consumers and in the second period when no breach was detected/announced respectively for $\lambda = 1$.

First, consider consumer surplus. For any given level of $q^f = q^c = q$,

$$\begin{aligned}CS_{mn}(q, q) - CS_{in}(q, q) &= \left(\frac{\hat{v}^n - \hat{v}^n_M}{2}\right)^2 + \delta(1 - (1 - q)\rho_B)\left(\frac{\hat{v}^n - \hat{v}^n_{ND}}{2}\right)^2 \\ &\geq 0.\end{aligned}$$

That is, consumer surplus is always higher when media notice is imposed for any given level of investment and belief. This implies that at $q = q^{in}_+$, the following inequality holds

$$CS_{mn}(q^*_{in}, q^*_{in}) \geq CS_{in}(q^*_{in}, q^*_{in}).$$

Since consumer surplus is increasing in q and $q^*_{mn} \geq q^*_{in}$, we have that

$$CS_{mn}(q^*_{mn}, q^*_{mn}) \geq CS_{mn}(q^*_{in}, q^*_{in}).$$

Combining the two inequalities, we obtain

$$CS_{mn}(q^*_{mn}, q^*_{mn}) \geq CS_{mn}(q^*_{in}, q^*_{in}) \geq CS_{in}(q^*_{in}, q^*_{in}).$$

Hence, we have shown that $CS_{mn}(q^*_{mn}, q^*_{mn}) \geq CS_{in}(q^*_{in}, q^*_{in})$.

I will now establish that the website's profit is also higher with media notice. By revealed preference,

$$\pi_{mn}(q^*_{mn}, q^*_{mn}) \geq \pi_{mn}(q^*_{in}, q^*_{mn}).$$

In addition, we have that

$$\begin{aligned}\pi_{mn}(q^*_{in}, q^*_{mn}) - \pi_{in}(q^*_{in}, q^*_{in}) &= (1 + \delta(1 - (1 - q^*_{in})\rho_B))(\hat{v}^n(q^*_{in} - \hat{v}^n(q^*_{mn}))r \\ &\geq 0.\end{aligned}$$

Hence,

$$\pi_{mn}(q^*_{in}, q^*_{mn}) \geq \pi_{in}(q^*_{in}, q^*_{in})$$

Together, the two conditions give

$$\pi_{mn}(q^*_{mn}, q^*_{mn}) \geq \pi_{mn}(q^*_{in}, q^*_{mn}) \geq \pi_{in}(q^*_{in}, q^*_{in}).$$

Thus, we have shown that profit is higher under the additional requirement of media notice.

Since both consumer surplus and profit are higher under media notice, social welfare is also higher.

Proof of Lemma 2

For any level of q, the bank is willing to invest at a given level of security γ whenever

$$t'(\gamma) \leq -\frac{\partial \phi_i(q, \gamma)}{\partial \gamma}, \quad i \in \{a, p\}.$$

Consider $\gamma \in [\overline{\gamma}, 1]$. Since t' is increasing in γ and $-\frac{\partial \phi_i(q,\gamma)}{\partial \gamma}$ is constant in this range, the value of $t'(\gamma) + \frac{\partial \phi_i(q,\gamma)}{\partial \gamma}$ is the smallest at $\gamma = \overline{\gamma}$. Thus, if $t'(\overline{\gamma}) + \frac{\partial \phi_i(q,\gamma)}{\partial \gamma}, > 0$, then for all $\gamma > \overline{\gamma}$,

$$t'(\gamma) > -\frac{\partial \phi_i(q, \gamma)}{\partial \gamma},$$

and the bank never chooses a level of security $\gamma > \overline{\gamma}$.

Further, for any set of parameter values, we know that the bank's marginal benefit of investment is decreasing in the website's level of investment q. Therefore, if

$$t'(\gamma) \leq -\frac{\partial \phi_i(1, \gamma)}{\partial \gamma}.$$

then,

$$t'(\gamma) > -\frac{\partial \phi_i(q, \gamma)}{\partial \gamma} \quad \forall q \in [0, 1].$$

Proof of Proposition 6

Existence

Let A denote the joint action space of the website and the bank; i.e. $A : [0,1] \times [0,1]$. Let f be the set of best response functions; i.e. $f : A \mapsto A$. Since A is a non-empty, compact and convex set and the best responses functions are continuous, by Brouwer's fixed point theorem, there exists a fixed point $x = (q_i^*, \gamma_i^*)$ such that $f(x) = x$. That is, a Nash equilibrium exists.

Uniqueness

Consider the interval $\gamma \in [0, \overline{\gamma}]$. For both "passive" and "active" security measures, it can be verified that γ_i is decreasing and convex in q_i. Further, $\gamma_i(1) = 0$ and $\gamma_i(0) < \overline{\gamma}$.

For "passive" measures, the website's best response function q_p is a decreasing and concave in γ_p. We know that $q_p(\gamma_p(1)) < 1$ and that $q_p(\gamma_p(0)) > q_p(\overline{\gamma}) > 0$, which means that the website's best response function lies below the bank's at $\gamma_p = 0$ and above it at $\gamma = \gamma_p(0)$. Since the two best response functions are continuous, they must cross at least once. Further, given that the best response functions are either strictly concave or strictly convex, the two curves cross only once and there is a unique equilibrium.

For "active" measures, the website's best response function is increasing and concave in γ. We have that, $q_a(\gamma_a(1)) = 0 < 1$ and $q_a(\gamma_a(0)) > 0$. Again, the continuity of the functions means that the functions must cross, and the strict concavity of the functions tells us that the functions cross only once. Hence, we have an unique equilibrium.

Proof of Proposition 7

Let k denote the probability that the consumer incurs a loss from using the website when bank's investment measures are "passive"; that is

$$k(q_p(\gamma_p), \gamma_p) = (1 - q_p(\gamma_p))\rho_B(1 - \gamma_p).$$

Consider a marginal increase in γ_p. This leads to an increase in the probability of loss whenever

$$\frac{dk(q_p(\gamma_p), \gamma_p)}{d\gamma_p} = -\rho_B(1 - q_p(\gamma_p) + (1 - \gamma_p)q_p'(\gamma_p)) > 0,$$

or equivalently,

$$-\frac{q_p'(\gamma_p)}{1 - q_p(\gamma_p)} > \frac{1}{1 - \gamma_p}. \tag{17}$$

It can be verified that the website's best-responsive function is decreasing and convex. Thus, we have that $-q'(\gamma_p^*) < -q'(\gamma_p)$ for all $\gamma_p \in [\gamma_p^*, \gamma_{min}]$. Further, we know that we know that the right hand side of the above expression is increasing in γ_p, which implies that $\frac{1-q_p(\gamma_{min})}{1-\gamma_{min}} > \frac{1-q_p(\gamma_p)}{1-\gamma_p}$ for all $\gamma \in [\gamma_p^*, \gamma_{min})$. Hence, raising the bank's security level from the unregulated equilibrium level to γ_{min} would necessarily results in an increase in the overall probability of loss when

$$-q_p'(\gamma_p^*) > \frac{1-q_p(\gamma_{min})}{1-\gamma_{min}}.$$

The slope of the website's best response function is

$$q_p'(\gamma_p) = -\frac{\delta\lambda\rho_B r}{c"(q_p(\gamma_p))}.$$

Substituting this into equation (6), we obtain:

$$c''(q_p^*) < \frac{\delta\lambda\rho_B r(1-\gamma_{min})}{1-q_p(\gamma_{min})},$$

which is the condition specified in the proposition.

Proof of Proposition 9

The full-crowding out of the website's investment exists whenever the bank's best response to the website's investment lies above the website's best response function in the interval $\gamma \in [0, \overline{\gamma}_m d]$. This implies that at $q_{md} = c'^{-1}(\delta\rho_B r)$, the optimal level of investment for the bank exceeds $\overline{\gamma}_{md}$:

$$\begin{aligned} \gamma_{md}(c'^{-1}(\delta\rho_B r)) &= t'^{-1}((1-c'^{-1}(\delta\rho_B r))\rho_B(1+\delta-\delta\rho_B)\alpha l) \\ &> \overline{\gamma}_{md} \\ \Leftrightarrow \quad t'(\overline{\gamma}_{md}) &< (1-c'^{-1}(\delta\rho_B r))\rho_B(1+\delta-\delta\rho_B)\alpha l. \end{aligned}$$

This gives us the condition stated in the proposition.