WORKING PAPERS

N° TSE-519

August 2014

"Ex Ante and Ex Post Investments in Cybersecurity"

Wing Man Wynne Lam

Toulouse School of Economics

Ex Ante and Ex Post Investments in Cybersecurity[*]

Wing Man Wynne Lam[†]

August, 2014

Abstract

This paper develops a theory of sequential investments in cybersecurity in which the software vendor can invest *ex ante* and *ex post.* The regulator can use safety standards and liability rules as means of increasing security. A standard is a minimum level of safety, and a liability rule states the amount of damage each party is liable for. I show that the joint use of an optimal standard and a full liability rule leads to underinvestment *ex ante* and overinvestment *ex post* because the software vendor does not suffer the full costs of the society in case of security failure. Instead, switching to a partial liability rule can correct the inefficiencies. This suggests that to improve security, the regulator should encourage not only the firms, but also the enterprises to invest in security. I also discuss the effect of network externality and explain why firms engage in "vaporware".

Keywords: cybersecurity, sequential investment, standards, liability

JEL Classification: L1, L8

1 Introduction

New security concerns are constantly arising as privacy breaches proliferate and cyber attacks escalate. For example, a recent data breach on an unprecedented scale saw more than 1.2 billion credentials stolen by a Russian criminal group.[1] Moreover, we continue to see the rise of "ransomware" (a malicious program that encrypts files on the victim's computer and demands a fee before unlocking those files), the discovery of security flaws on smartphones, and the emergence of new security risks from the "Internet of Things" (such as hackers stealing sensitive data from owners of Internet-connected objects—from locks, lights, thermostats, televisions, refrigerators, washing machines, to cars). A critical gap has thus emerged between firms' investment in cybersecurity and today's rapidly evolving technological advances, which warrants further research. More particularly, good security depends on more than just the

[*] I thank Giacomo Calzolari and Jacques Crémer for their helpful comments. Any opinions expressed are those of the author only.

[†] Toulouse School of Economics, University of Bologna. E-mail: wingmanwynne.lam2@unibo.it

[1] See "Russia gang hacks 1.2 billion usernames and passwords," *BBC News*, August 6 2014, available at `http://www.bbc.com/news/technology-28654613`.

technology. It requires a deeper understanding of the incentives of the agents who sell as well as those who use the technology. In the software industry, the incentives of those who are responsible for security and those who suffer from a security problem are often misaligned: while software vendors are motivated to minimize their own private costs, the social planner's goal is to minimize society's costs. Firms' incentives to invest are therefore suboptimal.[2]

The purpose of this paper is to understand how to use legislation such as safety standards and liability rules to provide incentives for software firms to make their product more secure. A standard is a minimum level of safety set by the regulator, and a liability rule states the amount of damage each party is liable for. In practice, there are different types of security standards, such as encryption standards, security breach notification standards, IT continuity standards, set by the National Institute of Standards and Technology (NIST) and Center for Internet Security (CIS) in the U.S., and more widely by the International Organization for Standards (ISO) and Internet Engineering Task Force (IETF). As for negligence liability, consumers continue to file lawsuits against firms for security breaches, data leakage, and infringement of privacy, and in this regard, these firms might be held accountable for consumer damages. This raises a number of interesting questions: Which of the interventions, standards or liability rules, would better incentivize firms and consumers to behave optimally? Should standards and liability rules be used separately or jointly? Is it socially optimal to shift some of the cost of investing in security from the firms to the consumers? To address these questions, I develop a model to study the investment incentives of a software firm when its software is subject to security problems and when consumers bear some precaution costs.

This paper makes two contributions. First, it studies a new type of inefficiency in the cybersecurity market, which is due to software vendors failing to take into account of consumers' cost of investing in security. Taking precautions is in general less costly for ordinary consumers as they only need to reboot their machines and the process of updating security is mostly automatic nowadays. However, the cost of precautions is significant for enterprise users, especially when they adopt sophisticated firewalls, cryptographic protocols, virus detection techniques, intrusion detection systems, data-loss prevention features, among others. Top-notch security tools are expensive and require a large number of man-hours to maintain and manage them. They are especially important for financial services, telecommunication sectors and government departments. Second, I introduce two types of investment the firms can undertake: *ex ante* care and *ex post* maintenance. In the software industry, as software is always evolving and adding new functionalities, they are never free of bugs. There are usually multiple rounds of debugging. Therefore, it is common for the software industry to have sequential investments. I further show that such possibility of sequential investments may lead to "vaporware" practice even in the absence of preemptive motives and reputation concerns: because *ex ante* and *ex post* investments are substitutes, allowing firms to identify security problem *ex post* increases the likelihood of releasing a less secure software product *ex ante*—a new perspective in the vaporware literature. In Sections 3.1 and 4, I also explore the consequences of public policies such as subsidizing the training of computer experts, synchronizing patch release and adoption

[2] See Anderson, Clayton and Moore (2009), and Anderson and Moore (2009) for surveys of the economics of network security.

cycles, and implementing vulnerability management by a third party.

To be more specific, I consider a model in which a firm sells software that is subject to potential security problems. The firm can invest *ex ante* to increase the security level and *ex post* to find the security problem before the hacker. If the firm discovers the bug, it can choose whether to disclose it or hide it. If the firm discloses the bug information, consumers can choose whether to take precaution or not. Consumers differ in their costs of taking precaution: actions are more costly for the laymen than for the computer experts.

I find that since the firm does not suffer the full costs of the society in case of security failure, its incentives to invest are suboptimal: it underinvests *ex ante* and overinvests *ex post.* I also show that there are inefficiencies associated with the joint use of a full liability rule and an optimal standard to increase security. Interestingly, a partial liability rule, which shifts some liability to the consumers, can correct the inefficiencies. This suggests that policies that encourage the consumers and the firm to share the costs of security could improve security. For example, since applying patches and malware-removal tools are costly for enterprise customers, the government could try to encourage them to put more effort in finding, testing and installing these tools as soon as the vendor makes them available. These results continue to hold in the presence of network externality.

I also show that if the firm has limited liability, increasing the number of computer experts mitigates suboptimal investment incentives. The reason is that the difference between the private and social incentives to invest arises from two effects. First, the firm does not pay fully for the damage, and the total amount of damage is decreasing in the number of experts. Second, the firm ignores the precautionary costs of the consumers when it makes its investment decision, and the total cost of precaution is increasing in the number of experts. When the firm has limited liability, the first effect dominates. This implies that to alleviate the inefficiency, the government can either impose limited liability on the firm and increase the number of computer experts, or simply allocate more liability to the firm. More particularly, under limited liability, the government can provide a subsidy for training in the area of cybersecurity so that enterprises become more competent in managing security threats. In contrast, if the firm bears substantial liability for consumers' damage, then the government needs to be careful about increasing the number of experts because the objectives of the planner and the firm will become more divergent.

1.1 Literature

This paper is primarily related to recent works on the economics of security investment. Gordon and Loeb (2002) and Kunreuther and Heal (2003) study the optimal security investment. Kunreuther and Heal (2003) consider the presence of network externality, but Gordon and Loeb (2002) do not. Both of them consider simultaneous investment, while I focus on sequential investment. Varian (2004) examines full liability in a model in which efforts of multiple parties are needed to increase security. He finds that liability should be assigned entirely to the party who can best manage the risk. Different from his analysis, I also consider partial liability, and the joint effect of partial liability and standards.

This paper also relates to the economics and legal literature on tort laws, but it departs from this literature by considering the possibility of consumers taking actions and sequential investments. More specifically, Shavell (1984) and Kolstad et al. (1990) compare standards with liability rules. However, Shavell's analysis is based on the inefficiencies associated with the potential bankruptcy of the firm and the uncertainty of lawsuit by the consumers, while the inefficiencies studied by Kolstad et al. are due to the uncertainty over the legal standard to which the firm will be held liable. Differently, inefficiencies here are caused by the firm failing to take into account of consumers' costs of investing in security. Moreover, the literature on torts has tended to focus on either *ex ante* investment, as in Daughety and Reinganum (1995, 2006), or *ex post* investment, as in Polinsky and Shavell (2010);[3][4] whereas this paper deals with both.

Finally, this paper shares with the literature on disclosure laws (see, for example, Granick (2005) and Choi et al. (2010)) the focus on the tradeoff that arises from disclosing software vulnerabilities: while secrecy prevents attackers from taking advantage of publicized security flaws, it interferes with scientific advancement in security, which is largely based on information sharing and cooperation. Choi et al. also examine the effect of a mandatory disclosure policy and a "bug bounty" program on welfare. However, they take security investments as given, and do not discuss optimal investment. Daughety and Reinganum (2005) study the effect of confidential settlement on product safety, but their focus is not on investment. This paper extends this literature by analyzing the optimal investment in security, and such investment is of two kinds: *ex ante* care and *ex post* maintenance.

2 The Model

Monopoly software vendor. Consider a firm that produces a software product which contains potential bugs. For simplicity, I assume away prices, so that the problem is simplified to choosing a level of security that minimizes the sum of the costs. The assumption is reasonable for consumers who have already bought the software and are therefore not concerned about the prices. Moreover, if the firm generates profit from channels other than selling the software product such as advertisement, then the objective is simply to minimize the costs.

Heterogeneous consumers. There is a unit mass of consumers. Consumers have different precaution costs: a proportion α of them are "computer experts" and have precaution cost γ drawn from a distribution $F(\gamma) \sim [0, +\infty)$, while the others are "laymen" with $\gamma = \infty$. Experts are security professionals who can take security precautions such as monitoring the system for attacks and patching the system if the firm discloses the presence of a security problem, while laymen without such professional knowledge will never take precautions.[5] In the main text,

[3] See Daughety and Reinganum (2013) for a survey of the literature on torts.

[4] Polinsky and Shavell analyze information acquisition about product risks when product quality is uncertain. Therefore, their problem concerns *ex post*, rather than *ex ante*, investment.

[5] I assume that consumers take precaution after the firm has disclosed the information about the bug. One could alternatively think of consumers taking precaution *ex ante*. However, the qualitative result will not change as long as the costs associated with these precautions are not borne by the firm.

all experts have the same γ and there are two types of consumers, but in Appendix A I show that the results are robust to the introduction of a continuum of consumer types. Assume that consumers always have positive utility in using the software.

Timing of the game. (i) The firm invests s in security at a cost $c(s)$. This is *ex ante* care. Such investment could take the form of improvement in infiltration detection or authentication technologies. (ii) By investing $m(b)$ in *ex post* maintenance, the firm will find a bug before the hacker does with probability b. Let $p(s)$ be the probability that the hacker will attack. I assume away strategic attacks.[6] (iii) If the firm discovers a bug, it can choose whether or not to disclose the security problem. Assume that there is no cost in disclosing the bug. For example, the firm can simply post the information on its website. However, disclosure increases the probability of attack by a small ϵ.[7] (iv) If the firm discloses a bug, the experts can choose whether or not to take precaution.

Assumption 1. *$c'(0) = 0, c'(s) > 0, c''(s) > 0, c'''(s) > 0, m'(0) = 0, m'(b) > 0, m''(b) > 0, m'''(b) > 0, p'(s) < 0,$ and $p''(s) > 0$.*

Under Assumption 1, investment costs $c(s)$ and $m(b)$ are thrice differentiable, convex, and increasing in s and b respectively;[8] and that probability of attack $p(s)$ is convex and decreasing in s.

Damage. For the firm, the damage incurred from an attack is $\overline{\eta}$ in case the hacker discovers the bug before the firm does, and $\underline{\eta}$ in case the firm identifies the bug first. Assume that $\overline{\eta} > \underline{\eta}$. This could be the financial loss caused by stolen information of the firm becoming available to the hacker. Such loss is smaller if the firm finds the bug first as it can then try to fix the problem. However, the firm may face substantial loss if the hacker exploits a bug that has not been previously identified—a phenomenon known as "zero-day attacks". For the consumers, the damage from an attack is $\overline{\mu}$ if they do not take precaution and $\underline{\mu}$ if they do. This could be monetary loss due to fraudulent use of their personal information. Assume that $\overline{\mu} > \underline{\mu}$, meaning once informed, consumers can take actions to mitigate the risk of being attacked. Let $\lambda \in [0, 1]$ denote the part of consumers' damages for which the firm is liable. I focus on three liability regimes:

- Full liability, under which the firm is liable for all damages faced by the consumers, i.e. $\lambda = 1$;
- Partial liability, under which the firm only compensates consumers partially, i.e. $\lambda \in (0, 1)$;
- No liability, under which consumers will not receive any compensation from the firm, i.e. $\lambda = 0$.

Thus, the total loss for the firm is $\eta + \lambda\mu$, where $\eta \in [\overline{\eta}, \underline{\eta}]$ and $\mu \in [\overline{\mu}, \underline{\mu}]$.

[6] Strategic attacks are modeled in, for instance, Acemoglu et al. (2013). They show that strategic targeting provides additional incentives for overinvestment in security because larger investment shifts attacks from one agent to another.

[7] Arora, Nandkumar and Telang (2006) show empirically that in some cases vulnerability disclosure increases the frequency of attacks.

[8] The third derivatives ensure that the profit function is well-behaved.

3 Optimal Investment

I now work backward from the last stage. When the firm discloses a bug, the expected damage for a consumer who does not take precaution is $p(s)\overline{\mu}$, and that for a consumer who takes precaution is $p(s)\underline{\mu} + \gamma$. Therefore, the consumer will take precaution if

$$\gamma < p(s)(\overline{\mu} - \underline{\mu}). \tag{1}$$

In the disclosure stage, the firm can choose its disclosure policy in case it discovers a bug. If it does not disclose the security problem, its expected cost is $p(s)(\underline{\eta} + \lambda\overline{\mu})$. If it chooses to disclose, there are two cases. If consumers take precaution, the firm incurs a cost of $p(s)[\underline{\eta} + \lambda(\alpha\underline{\mu} + (1-\alpha)\overline{\mu})]$. However, if consumers do not take precaution, the cost becomes $p(s)(\underline{\eta} + \lambda\overline{\mu})$.[9] Therefore, the firm will only disclose if this leads consumers to take precaution, that is, if Equation (1) holds.

In the investment stage, the firm chooses s and b to minimize its expected loss, which is denoted by $\mathcal{L}^f$.

$$\begin{aligned}\min_{b,s} \mathcal{L}^f = {} & (1-b)p(s)(\overline{\eta} + \lambda\overline{\mu}) \\ & + b\left\{\int_0^{p(s)(\overline{\mu}-\underline{\mu})} p(s)[\underline{\eta} + \lambda(\alpha\underline{\mu} + (1-\alpha)\overline{\mu})]dF(\gamma) + \int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} p(s)(\underline{\eta} + \lambda\overline{\mu})dF(\gamma)\right\} \\ & + m(b) + c(s). \end{aligned} \tag{2}$$

Let $b^m(s)$ denote the firm's optimal *ex post* investment strategy given *ex ante* security s, and let s^* and $b^* \equiv b^m(s^*)$ denote the solutions of Equation (2).

The first term in Equation (2) is the expected cost of the firm when the hacker discovers the bug first, and in which case both the firm and the consumers suffer a large damage. When the firm finds the bug before the hacker, either it discloses the bug if consumers' cost is small, which is captured by the second term, or it does not disclose if consumers' cost is large, which is captured by the third term. In this case, the firm suffers a small damage from attack because it identifies the bug sooner than the hacker, while the extent of damages suffered by the consumers depends on whether precautionary measures are taken. The last two terms represent *ex ante* and *ex post* investment costs.

The social planner's incentive to disclose is the same as the firm, that is, the planner will disclose as long as γ is small enough. However, different from the firm, if the planner chooses to disclose, its expected cost is $p(s)(\underline{\eta} + \alpha\underline{\mu} + (1-\alpha)\overline{\mu}) + \alpha\gamma$, which is higher than that of the firm. This is because the planner also takes into account consumers' cost of taking precautions, and internalizes all the costs, so there is no liability issue. In case of non-disclosure, the expected cost is $p(s)(\underline{\eta} + \overline{\mu})$.

The social planner chooses s and b to minimize the expected loss of the society, which is

[9] When consumers do not take precaution, the firm is indifferent between disclosing and not disclosing. However, by assuming that disclosure would increase the probability of attack by ϵ, the firm will strictly prefer not to disclose.

denoted by $\mathcal{L}^{SP}$.

$$\begin{aligned}\min_{b,s}\mathcal{L}^{SP} =&(1-b)p(s)(\overline{\eta}+\overline{\mu})+b\left\{\int_0^{p(s)(\overline{\mu}-\underline{\mu})}[p(s)(\underline{\eta}+\alpha\underline{\mu}+(1-\alpha)\overline{\mu})+\alpha\gamma]dF(\gamma)\right.\\&\left.+\int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty}p(s)(\underline{\eta}+\overline{\mu})dF(\gamma)\right\}+m(b)+c(s)\\=&\mathcal{L}^f|_{\lambda=1}+b\alpha\int_0^{p(s)(\overline{\mu}-\underline{\mu})}\gamma dF(\gamma).\end{aligned} \tag{3}$$

Let $b^{SP}(s)$ denote the social planner's optimal *ex post* investment strategy given *ex ante* security s, and let s^o and $b^o \equiv b^{SP}(s^o)$ denote the solutions of Equation (3).

The difference between $\mathcal{L}^f$ and $\mathcal{L}^{SP}$ is that the firm minimizes its own private costs, while the social planner minimizes the sum of firm's and consumers' costs.

Lemma 1. *Under full liability* ($\lambda = 1$), $b^m(s)$ *and* $b^{SP}(s)$ *decrease with* s.

Proof. See Appendix B. □

Lemma 1 shows that the firm has less incentive to find bugs *ex post* given a high security level *ex ante*, meaning that *ex ante* and *ex post* investments are substitutes.

Lemma 2. *Under full liability* ($\lambda = 1$), $b^m(s) > b^{SP}(s)$ *for all* s. *In particular, if the standard is set at the socially optimal level,* $s^* = s^o$, *the firm will overinvest in ex post maintenance,* $b^m(s^o) > b^{SP}(s^o)$.

Proof. See Appendix C. □

One might expect that under full liability and an optimal standard the firm will invest optimally, but it turns out differently when consumers also bear some costs in protecting their computers. The intuition runs as follows. If a bug is not found, both the firm and the society suffer the same loss. If a bug is discovered, the firm can reduce the loss more than the planner because it does not bear the costs of the consumers. Since the firm has more to gain in finding the bug, it will overinvest.

I assume that full liability is defined for "net" damages to the consumers. One can alternatively define it for "total" damages, which includes also consumers' precaution cost. In this case, full liability alone is enough to restore the first-best. I model the liability regime the way I did because in practice, firms are typically liable for financial damages to the consumers caused by, for example, a data breach. Liability sometimes also covers for litigation costs, but very rarely for investment costs in precaution. One difficulty lies in estimating the amount of time and effort consumers spent on managing, maintaining and patching a system.

Proposition 1. *(Full Liability). Under full liability* ($\lambda = 1$), *the firm underinvests in ex ante care,* $s^* < s^o$, *and overinvests in ex post maintenance,* $b^* > b^o$.

Proof. See Appendix D. □

Proposition 1 shows that full liability alone does not achieve the first-best solution. The reason is that, as shown in Lemma 2, *ex post* the firm has more to gain in finding the bug than the planner, and hence it invests too much in *ex post* maintenance. The firm invests too little in *ex ante* care because it expects to overinvest *ex post*, as was shown in Lemma 1.

Proposition 2. *(Partial Liability). The socially optimal level of investment, s^o and b^o, can be achieved with the joint use of an optimal standard s^o and a partial liability rule $\lambda \in (0,1)$.*

Proof. See Appendix E. □

When security standards are set at the socially optimal level, it is inefficient to implement full liability because the firm will overinvest *ex post*; it is also inefficient to set firm's liability to zero because it will then underinvest *ex post*. As a consequence, the optimal liability rule is a partial one. Note that in Appendix F I show that if liability regime is the only instrument of public policies, it is not enough to provide the right incentives for two investments.

3.1 Network Externality

In this subsection, I consider direct and indirect network effects. In practice, users whose computers are infected may create negative externalites on the other users in that attackers can use these computers to host phishing sites, distribute spam e-mails or other unlawful content. Kunreuther and Heal (2003), August and Tunca (2006), Acemoglu et al. (2013), and Riordan (2014), for instance, examine agents' incentive to invest in security under the presence of network externalities. While they focus on one type of security investment, this paper deals with two types.[10]

Let us first examine the situation with indirect network effects in which the firm's investment strategy is affected by the proportion of consumers taking precaution.

Corollary 1. *(Indirect network effects). When λ is large, increasing the proportion of computer experts, α, exacerbates the ex ante underinvestment and ex post overinvestment problems. When λ is small, increasing α mitigates the investment problem.*

Proof. See Appendix G. □

The intuition behind Corollary 1 runs as follows. Comparing Equations (2) with (3), the difference between the private and social incentives to invest that is related to α arises from the following.

$$p(s) \quad \underbrace{(1-\lambda)(\alpha\underline{\mu} + (1-\alpha)\overline{\mu})}_{distortion\ from\ liability\ assignment} \quad + \quad \underbrace{\alpha\gamma}_{distortion\ from\ consumers'\ costs} \quad .$$

[10] More particularly, August and Tunca (2006) focus on the problem of patch management, and therefore consider *ex post* investment only. Security investments are strategic complements in Kunreuther and Heal (2003), strategic substitutes in Acemoglu et al. (2013), and can be strategic complements or strategic substitutes in Riordan (2014) depending on whether the attacks are direct or indirect, but agents can only invest once in these models.

Investment incentives are therefore distorted by two forces: first, the firm does not pay fully for the damage; second, the firm ignores the precautionary costs of the consumers when it makes its investment decision. If the firm is held liable for a large proportion of damage (i.e. λ is large), then reducing the proportion of experts (α) mitigates suboptimal investment incentives. The reason is that an increase in firm's liability reduces the first type of distortion, whereas a decrease in the proportion of experts reduces the second type of distortion. Taking the effects together, the objectives of the planner and the firm become more aligned, and thus this reduces the extent that the firm is investing suboptimally. If, on the other hand, the firm is held liable for a smaller proportion of damage, then increasing the proportion of experts will reduce the inefficiency. This is because the extent of the first type of distortion depends on the total amount of damage, and is decreasing in α, whereas the extent of the second type of distortion depends on the total cost of precaution of the consumers, and is increasing in α. When the firm has limited liability, the first type of distortion dominates.

This implies that to alleviate the inefficiency, the government can either impose limited liability on the firm and increase the number of computer experts, or simply allocate more liability to the firm. More particularly, under limited liability, the government can provide a subsidy for training in the area of cybersecurity so that enterprises become more competent in managing security threats. For example, many security breaches involve attackers trying to compromise users' accounts, and users are sometimes unaware of such attack. Even if they are aware of the attack, they sometimes lack the skills needed to resolve the security problem. Therefore, increasing training that aims to enhance the technical skills of these enterprise users appears to be appropriate provided that the cost of implementing this subsidy is not too large. In contrast, if the firm bears substantial liability for consumers' damage, then the government needs to be careful about increasing the number of experts because the objectives of the planner and the firm would further diverge. That being said, this does not mean that offering cybersecurity training is undesirable (e.g. it could potentially generate cost savings for firms through detecting, defending against and recovering from cyber-attacks), but that the potential adverse effects on incentives should not be ignored.

Previously, I have assumed that there are no direct network effects, but my qualitative results would not change even if we add this. Re-interpreting *ex post* investment as a patch release and consumers' action as the choice of patch installation, direct network effects between consumers could arise when consumers who do not patch increase the security risks on other consumers, and consumers who patch reduce the probability of others being attacked. In this case, increasing the proportion of experts α will lower the damage to all experts, $\underline{\mu}$, and that to all laymen, $\overline{\mu}$, meaning only magnitude changes. However, the main qualitative result of liability-sharing between the firm and the consumers remains valid, provided consumers have to take precautionary actions.

4 Discussion

Vaporware.—"Vaporware" refers to the software industry practice of announcing new products well in advance of their actual release on the market.[11] The previous literature, for instance, Bayus et al. (2001) and Haan (2003), studies how such product pre-announcements can be used as a means of entry deterrence in a signaling model. Choi et al. (2010) examine how reputation concerns may induce firms to make honest announcements in a repeated cheap-talk game. Although vaporware practice typically means the release dates of the products are much later than the original announced dates, we could alternatively view the announced product as a product characteristics (a security feature, for instance) instead of the physical product. Vaporware could then be interpreted as delivering a lower-quality product than promised, which is consistent with the current development in the industry: software products, mobile applications, and smart-home appliances are often launched prematurely while they are still in development and are therefore susceptible to security risks. The result of *ex ante* underinvestment in security in this model captures the essence of this situation. Moreover, I show that underinvestment may occur even in the absence of preemptive motives and reputation concerns. This is therefore different from the vaporware literature, where firms engage in vaporware only to prevent entry or when reputational concern is not so important. The new insight here is that the possibility of sequential investments, which allows the firm also to invest *ex post* in fixing the security problem, provides an alternative explanation at least in part for vaporware practice in the software market.

Policy Implications.—I have examined the investment incentive of a software vendor, both *ex ante* and *ex post*, when consumers bear some costs of taking precaution. I find that security can be improved with the joint use of an optimal standard and a partial liability rule. This implies that the regulator can enforce some minimum standards for encryption and security breach notification. Sanctions can be imposed if these requirements are violated. Another policy we can consider is liability regime. Interestingly, I find that, given an optimal standard, shifting some liability to the consumers is welfare improving. This means that the regulator should not impose a one hundred percent liability on the software vendor because this will distort its investment incentives. Instead, an effective policy is to ask both the software vendor and its customers to share the costs of security.[12]

Despite the fact that users dislike or feel concerned about security problems, some of them ignore notifications from the vendor and do not take up any of the proposed solutions. For example, more than 90% of ChoicePoint customers whose personal information had been stolen did not take up the mitigating solutions offered by the firm such as free credit monitoring service and insurance after the data breach.[13] This may be due to the fact that consumers have other

[11] Vaporware may also mean the announced products never reach the market, but this is not the focus of this paper because the firm always introduces the product in this model.

[12] Although this discussion interprets costs of security as a form of liability, they are different from the costs explained by γ in that consumers ignoring or not noticing security alerts is not an investment, but rather it shows a systematic lack of security consciousness. This raises the question of who should be responsible for the damages that arise from such negligence.

[13] See Jon Brodkin, "Victims of ChoicePoint Data Breach Didn't Take Advantage of Free Offers," *Network*

competing demands on their time, and paying attention to data breach notifications appears to be low on their priority list.

On enterprise level, installing patches could be costly especially for large companies because the plethora of security updates can often overwhelm software engineers, who have to keep track of all relevant bugs and patches, and match the version of all those updates to the version of software their company is using. Once a problem is identified, they need to figure out which updates get priority, and look for solutions to deal with it.[14] In addition, if the installation requires rebooting an enterprise's critical system, downtime can be expensive. As a consequence, this could easily lead to the missing of some major security updates.

This suggests that a desirable policy should try to eliminate the delay in applying the solutions to security problems. First, the government could persuade or mandate the users to react more quickly (for example, within a predetermined window of time) as soon as the vendor makes the solutions available and notifies them in a reasonable way. Second, third parties can be introduced to help enterprises to find, select and deploy the solutions that are relevant to their systems. An example of third-party vulnerability management that helps businesses to adhere to compliance and security standards in the IT and financial sectors is Qualys, Inc.

More General Applications.—The analysis also provides insight into other industries in which sequential investments are important, such as automobiles and pharmaceuticals. We can then re-interpret the seller as a firm that produces a product with some safety features. There are again two types of investments the firm can undertake: *ex ante* investment in pre-sale product design, and *ex post* investment in post-sale product testing. For example, *ex ante* investment could lead to the development of a new technology in cars that is subject to potential safety defect, or a new drug that has previously unknown side effects. The firm can invest *ex post* to remedy these safety problems. We can then use the previous analysis to study investment incentives of the firm, in particular whether there are incorrect incentives to provide safety *ex ante* and *ex post* and how to improve them.

5 Conclusion

To increase security, the key is not so much about holding the software vendor solely liable for the loss, but balancing the investment incentives between different players. This discussion represents a useful first step towards understanding sequential security investments. In future work, it might be interesting to relax the single-firm assumption and consider dynamic issues and contagion issues in a network of multiple firms.[15]

World, April 10, 2007, http://www.networkworld.com/news/2007/041007-choicepoint-victim-offers.html?page=1.

[14] Practitioners have commonly considered patch management as a time- and resource-consuming activity. See, for instance, Symantec, "Automating Patch Management," February 8, 2005, http://www.symantec.com/articles/article.jsp?aid=automating_patch_management.

[15] See, for instance, Morris (2000), Acemoglu et al. (2013), and Goyal et al. (2014) for treatment of contagion in networks.

Appendices

A Continuum of Consumers

With a slight abuse of the notation, suppose that there is a continuum of consumers whose precaution cost γ is drawn from a distribution $F(\gamma) \sim [0, +\infty)$. As before, consumers will take precaution if $\gamma < p(s)(\overline{\mu} - \underline{\mu})$, and the marginal consumer, who is indifferent between taking and not taking precaution, is given by $\gamma(s) \equiv p(s)(\overline{\mu} - \underline{\mu})$.

If the firm does not disclose the bug, its expected cost is $p(s)(\underline{\eta} + \lambda\overline{\mu})$; if it discloses the bug, it expected cost is $p(s)[\underline{\eta} + \lambda(F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu})]$. Since the latter is smaller than the former, the firm will always disclose. Therefore, the firm chooses s and b to minimize

$$\min_{b,s} \mathcal{L}^f = (1-b)p(s)(\overline{\eta} + \lambda\overline{\mu}) + bp(s)[\underline{\eta} + \lambda(F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu})] + m(b) + c(s). \quad \text{(A.1)}$$

As for the planner, the cost for non-disclosure is $p(s)(\underline{\eta} + \overline{\mu})$, whereas the cost for disclosure is $p(s)[\underline{\eta} + F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu}] + \int_0^{\gamma(s)} \gamma dF(\gamma)$. Since the latter is smaller than the former, the planner will always disclose. The planner therefore solves

$$\begin{aligned} \min_{b,s} \mathcal{L}^{SP} &= (1-b)p(s)(\overline{\eta} + \overline{\mu}) \\ &\quad + b\left\{ p(s)[\underline{\eta} + F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu}] + \int_0^{\gamma(s)} \gamma dF(\gamma) \right\} + m(b) + c(s). \quad \text{(A.2)} \end{aligned}$$

It is easy to see that since $\int_0^{\gamma(s)} \gamma dF(\gamma) > 0$, $\mathcal{L}^{SP} > \mathcal{L}^f$ for any λ. Thus, the main results of *ex ante* underinvestment and *ex post* overinvestment carry through.

B Proof of Lemma 1

Since $\lambda = 1$, the first-order conditions with respect to b are given by

$$\begin{aligned} \frac{\partial \mathcal{L}^{SP}}{\partial b} &= 0, \\ \Leftrightarrow m'(b) &= p(s)(\overline{\eta} + \overline{\mu}) - \underbrace{\int_0^{p(s)(\overline{\mu} - \underline{\mu})} [p(s)(\underline{\eta} + \alpha\underline{\mu} + (1-\alpha)\overline{\mu}) + \alpha\gamma] dF(\gamma)}_{G^{SP}(s)} \\ &\quad - \int_{p(s)(\overline{\mu} - \underline{\mu})}^{\infty} p(s)(\underline{\eta} + \overline{\mu}) dF(\gamma), \quad \text{(B.1)} \end{aligned}$$

and

$$\begin{aligned} \frac{\partial \mathcal{L}^{f}}{\partial b} &= 0, \\ \Leftrightarrow m'(b) &= p(s)(\overline{\eta} + \overline{\mu}) - \underbrace{\int_0^{p(s)(\overline{\mu} - \underline{\mu})} p(s)(\underline{\eta} + \alpha\underline{\mu} + (1-\alpha)\overline{\mu}) dF(\gamma)}_{G^{f}(s)} \\ &\quad - \int_{p(s)(\overline{\mu} - \underline{\mu})}^{\infty} p(s)(\underline{\eta} + \overline{\mu}) dF(\gamma). \quad \text{(B.2)} \end{aligned}$$

The right hand sides of Equations (B.1) and (B.2) are decreasing in s.

C Proof of Lemma 2

We can see from Equations (B.1) and (B.2) that if $s^* = s^o$, then $G^f(s^o) < G^{SP}(s^o)$. Thus, $b^m(s^o) > b^{SP}(s^o)$.

D Proof of Proposition 1

Since $\lambda = 1$, the first-order conditions with respect to s are given by

$$\begin{aligned} \frac{\partial \mathcal{L}^{SP}}{\partial s} &= 0, \\ \Leftrightarrow -\frac{c'(s)}{p'(s)} &= (1-b)(\overline{\eta}+\overline{\mu}) + b\left[\int_0^{p(s)(\overline{\mu}-\underline{\mu})} (\underline{\eta}+\alpha\underline{\mu}+(1-\alpha)\overline{\mu})dF(\gamma)\right. \\ &\quad \left. + \int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} (\underline{\eta}+\overline{\mu})dF(\gamma)\right], \end{aligned} \tag{D.1}$$

and

$$\begin{aligned} \frac{\partial \mathcal{L}^{f}}{\partial s} &= 0, \\ \Leftrightarrow -\frac{c'(s)}{p'(s)} &= (1-b)(\overline{\eta}+\overline{\mu}) + b\left[\int_0^{p(s)(\overline{\mu}-\underline{\mu})} (\underline{\eta}+\alpha\underline{\mu}+(1-\alpha)\overline{\mu})dF(\gamma)\right. \\ &\quad \left. + \int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} (\underline{\eta}+\overline{\mu})dF(\gamma) - \alpha p(s)(\overline{\mu}-\underline{\mu})^2 f(p(s)(\overline{\mu}-\underline{\mu}))\right]. \end{aligned} \tag{D.2}$$

Define the right hand side of Equation (D.1) as $H^{SP}(b)$, and that of Equation (D.2) as $H^f(b)$. Clearly, the left hand sides of Equations (D.1) and (D.2) are equal. However, $H^{SP}(b^{SP}(s)) > H^f(b^{SP}(s)) > H^f(b^m(s))$. The first inequality follows from $H^{SP}(b) > H^f(b)$ for any b, whereas the second inequality is due to the fact that $H^f(b)$ is decreasing in b.

Since $c'''(s) > 0$ and $p'''(s) > 0$, it is easy to see that $-c'(s)/p'(s)$ is convex and increasing in s, and it has the limits $\lim_{s\to 0} -c'(s)/p'(s) = 0$ and $\lim_{s\to\infty} -c'(s)/p'(s) = \infty$. As for the right hand sides, the limits of both $H^{SP}(b)$ and $H^f(b)$ are bounded away from ∞ as s tends to ∞. Moreover, $H^{SP}(0) > 0$, and if $H^f(0) > 0$, the solution to both equations exists, and we denote them by s^* and s^o respectively. In addition, if the solution is unique, we must have $s^* < s^o$ due to the fact that $H^{SP}(b^{SP}(s)) > H^f(b^m(s))$.[16]

Using Lemma 1, if $s^* < s^o$, then $b^* > b^o$.

[16] For example, there exists a unique equilibrium investment when both $F(p(s))$ and $p(s)f(p(s))$ are convex, and $m(b)$ is quadratic.

E Proof of Proposition 2

Suppose $s^* = s^o$. If $\lambda = 1$, Lemma 2 implies $b^m(s^o) > b^{SP}(s^o)$. If $\lambda = 0$, Equation (B.2) becomes

$$m'(b) = p(s)(\overline{\eta} - \underline{\eta}).$$

Comparing with Equation (B.1), $b^m(s^o) < b^{SP}(s^o)$. Therefore, there exists $\lambda \in (0, 1)$ such that $b^m(s^o) = b^{SP}(s^o)$.

F Liability regime as the only instrument

Suppose that there exists $\lambda \in [0, 1]$ such that $b^* = b^o$ and $s^* = s^o$. This implies that $\partial \mathcal{L}^f / \partial b = \partial \mathcal{L}^{SP} / \partial b$ and $\partial \mathcal{L}^f / \partial s = \partial \mathcal{L}^{SP} / \partial s$. However, we can easily verify that these two conditions cannot be satisfied at the same time.

G Proof of Corollary 1

The difference between Equations (B.1) and (B.2) is

$$m'(b^*) - m'(b^o) = \alpha \int_0^{p(s)(\overline{\mu}-\underline{\mu})} \gamma dF(\gamma),$$

which is positive and increasing in α, meaning that a larger α worsens the *ex post* overinvestment problem.

Similarly, the difference between Equations (D.1) and (D.2) is

$$(b^* - b^o)\left[\int_0^{p(s)(\overline{\mu}-\underline{\mu})} (\underline{\eta} + \alpha\underline{\mu} + (1-\alpha)\overline{\mu})dF(\gamma) + \int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} (\underline{\eta} + \overline{\mu})dF(\gamma) - (\overline{\eta} + \overline{\mu})\right]$$
$$- \alpha b^* p(s)(\overline{\mu} - \underline{\mu})^2 f(p(s)(\overline{\mu} - \underline{\mu})).$$

The first term $(b^* - b^o)$ is positive and increasing in α, and the term in the square bracket is negative and decreasing in α. The product of these two terms is thus negative and decreasing α. Since the final term $-\alpha b^* p(s)(\overline{\mu} - \underline{\mu})^2 f(p(s)(\overline{\mu} - \underline{\mu}))$ is also negative and decreasing in α, taken together the difference between Equations (D.1) and (D.2) is negative and decreasing in α, meaning that the *ex ante* underinvestment problem is more severe as α increases.

This proof remains valid as long as λ is large enough.

References

[1] Daron Acemoglu, Azarakhsh Malekian, and Asu Ozdaglar. Network Security and Contagion. MIT Working Paper, 2013.

[2] Ross Anderson, Richard Clayton, and Tyler Moore. The Economics of Online Crime. *Journal of Economic Perspectives*, 23(3):3–20, 2009.

[3] Ross Anderson and Tyler Moore. Information Security: Where Computer Science, Economics and Psychology Meet. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 367(1898):2717–2727, 2009.

[4] Ashish Arora, Anand Nandkumar, and Rahul Telang. Does information security attack frequency increase with vulnerability disclosure? An empirical analysis. *Information Systems Frontiers*, 8(5):350–362, 2006.

[5] Terrence August and Tunay Tunca. Network Software Security and User Incentives. *Management Science*, 52(11):1703–1720, 2006.

[6] Barry Bayus, Sanjay Jain, and Ambar Rao. Truth or Consequences: An Analysis of Vaporware and New Product Announcements. *Journal of Marketing Research*, 38(1):3–13, 2001.

[7] Jay Pil Choi, Chaim Fershtman, and Neil Gandal. Network Security: Vulnerabilities and Disclosure Policy. *Journal of Industrial Economics*, 58(4):868–894, 2010.

[8] Jay Pil Choi, Eirik Kristiansen, and Jae Nahm. Vaporware. *International Economic Review*, 51(3):653–669, 2010.

[9] Andrew Daughety and Jennifer Reinganum. Product Safety: Liability, R&D and Signaling. *American Economic Review*, 85(5):1187–1206, 1995.

[10] Andrew Daughety and Jennifer Reinganum. Secrecy and Safety. *American Economic Review*, 95(4):1074–1091, 2005.

[11] Andrew Daughety and Jennifer Reinganum. Markets, Torts and Social Inefficiency. *RAND Journal of Economics*, 37(2):300–323, 2006.

[12] Andrew Daughety and Jennifer Reinganum. Economic Analysis of Products Liability: Theory. In Jennifer Arlen, editor, *Research Handbook on the Economics of Torts*, chapter 3, pages 69–96. Edward Elgar Publishing Ltd., 2013.

[13] Lawrence Gordon and Martin Loeb. The Economics of Information Security Investment. *ACM Transactions on Information and System Security*, 5(4):438–457, 2002.

[14] Sanjeev Goyal, Hoda Hiedari, and Michael Kearns. Competitive Contagion in Networks. *Games and Economic Behavior*, 2014, forthcoming.

[15] Jennifer Granick. The Price of Restricting Vulnerability Publications. *International Journal of Communications Law & Policy*, 9:1–35, 2005.

[16] Marco Haan. Vaporware as a Means of Entry Deterrence. *Journal of Industrial Economics*, 51(3):345–358, 2003.

[17] Charles Kolstad, Thomas Ulen, and Gary Johnson. Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements? *American Economic Review*, 80(4):888–901, 1990.

[18] Howard Kunreuther and Geoffrey Heal. Interdependent Security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, 2003.

[19] Stephen Morris. Contagion. *Review of Economic Studies*, 67(1):57–78, 2000.

[20] A. Mitchell Polinsky and Steven Shavell. Mandatory Versus Voluntary Disclosure of Product Risks. *Journal of Law, Economics, & Organization*, 28(2):360–379, 2010.

[21] Michael Riordan. Economic Incentives for Security. Powerpoint Slides presented at Cybercriminality Seminar at Toulouse School of Economics on 4 June, 2014.

[22] Steven Shavell. A Model of the Optimal Use of Liability and Safety Regulation. *RAND Journal of Economics*, 15(2):271–280, 1984.

[23] Hal Varian. System Reliability and Free Riding, 2004. Available at `http://people.ischool.berkeley.edu/~hal/Papers/2004/reliability` (accessed 1 December, 2013).