Attack-Prevention and Damage-Control Investments in Cybersecurity*

Wing Man Wynne Lam†

September, 2016

Abstract

This paper examines investments in cybersecurity made by users and software providers with a focus on the latter's concerning attack prevention and damage control. I show that full liability, whereby the provider is liable for all damage, is inefficient, owing namely to underinvestment in attack prevention and overinvestment in damage control. On the other hand, the joint use of an optimal standard, which establishes a minimum compliance framework, and partial liability can restore efficiency. Implications for cybersecurity regulation and software versioning are discussed.

Keywords: Cybersecurity; Investment; Standard; Liability; Bilateral care

JEL Classification: K13, L1, L8

*I thank the editor, Christiaan Hogendorn, and two anonymous referees for comments that significantly improved the paper. I also thank Paul Belleflamme, Giacomo Calzolari, Jacques Crémer, Vincenzo Denicolò, Axel Gautier, Domenico Menicucci, Paul Seabright, and the participants at the LCII and CORE 2015 Digital Economy Workshop, the WEIS 2015 Conference, and the EALE 2015 Conference, the George Washington University 2016 Cybersecurity Workshop, as well as those at seminars at Télécom ParisTech and Saint-Louis University Brussels for their valuable comments. I also acknowledge the support of Toulouse School of Economics and University of Bologna at earlier stages of this research. All opinions expressed are strictly my own.

†University of Liege (ULg), HEC Management School, Liege Competition and Innovation Institute (LCII). E-mail: wingmanwynne.lam@ulg.ac.be

1 Introduction

New security concerns are constantly arising as privacy breaches proliferate and cyber attacks escalate. For example, a recent data breach at Dropbox has affected more than 68 million users.[1] And, as persistent are 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 of the "Internet of Things" (such as hackers stealing sensitive data from owners of Internet-connected objects—from locks, lights, thermostats, televisions, refrigerators, and washing machines to cars). It is also common to see software providers releasing vulnerable alpha versions of their products before the more secure beta versions. Thus, a critical lag has emerged between software providers' investment in cybersecurity and today's rapidly evolving technological advances. This paper presents a model accounting for the investment incentives of the various parties affected by security concerns and analyzing the appropriate remediation when these incentives depart from the socially efficient level.

A prominent feature of the software industry is its fast development and release of new functionalities. Software products are therefore never free of bugs, and it is very common to observe multiple rounds of investments. To incorporate this feature, I consider a software provider that sells a software product—subject to potential security problems—and can invest in attack prevention and damage control to increase security. Attack-prevention investments, e.g., good infiltration detection and authentication technologies, reduce the probability of successful attacks (among others, phishing, denial-of-service, virus attacks). On the other hand, damage-control investments are remediation strategies, e.g., finding, testing, and fixing bugs reduces the probability that the hacker finds and exploits a bug before the provider does. Both types of investments are crucial to raising the security level of a product. For instance, Gartner predicts that both investments in attack prevention and damage control will continue to grow, as organizations focusing just on one have not been successful in increasing security.[2]

Software users can also invest in security. If the provider finds and discloses a bug, then users can adopt various defenses (among others, user-side encryption, firewalls, virus detection techniques, intrusion detection systems, data-loss prevention features) against online attacks. Not all users, however, whether enterprise and home users, take preventive measures even when software providers disclose bug information. Kaspersky Lab reports that hackers often use exploits for known vulnerabilities against enterprises because enterprises are slow to apply patches. For example, the use of exploits for office software vulnerabilities against enterprise users is three times as frequent as that against home users.[3] Symantec (2016) also reports that more than 75% of websites Symantec scanned contained unpatched vulnerabilities in 2015, with very little improvement over the past three years. We capture the lack of precautionary

[1] See "Dropbox hack affected 68 million users," *BBC News*, August 31, 2016, available at `http://www.bbc.com/news/technology-37232635`.

[2] See "Gartner says Worldwide Information Security Spending will grow 7.9% to reach $81.6 billion in 2016," *Gartner Press Release*, August 9, 2016, available at `http://www.gartner.com/newsroom/id/3404817`.

[3] See "Evolution of cyber threats in the corporate sector," *Kaspersky Security Bulletin 2015*, December 10, 2015, available at `http://bit.ly/2bQ4Q1C`.

actions by assuming that there are costs of taking precaution. Furthermore, depending on the magnitude of these costs, a user is either a layman or an expert: actions are more costly for the former than for the latter. For example, the costs of taking precautions vary between different types of enterprise users. While financial services, telecommunication sectors, utilities and government departments have far more resources to hire security professionals to maintain and manage top-notch security tools, smaller companies have relatively limited budgets to that effect. Hence, their engineers may not have a keen understanding about the state-of-the-art security, which results in higher learning costs than their more advanced counterparts. For short, there are three types of investments: the provider's attack-prevention investment reduces the attack occurring probability (once an attack occurs, it causes damage to both the provider and the users), whereas the provider's bug-fixing investment and user precautionary actions limit the extent of the damage.

To eliminate potential investment inefficiencies, the regulator can ideally impose the optimal levels of attack-prevention and damage-control investments whenever both investments are observable and verifiable. In reality, however, it is difficult to monitor a provider's investment in bug discovery and bug fixing: because the objective is to find hitherto unknown vulnerabilities, the success of discovery, which depends on rapidly evolving attack and defense technologies, is largely uncertain. On the contrary, attack-prevention investment is relatively easy to monitor as its objective is to defend against known vulnerabilities. For instance, new software products can be tested for known vulnerabilities to ensure that they are secure before they can be released on the market. Thus, I assume that the regulator can regulate directly attack-prevention—but not damage-control—investment by setting a standard (i.e., a minimum level of security). 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). Similarly, in the banking industry, the Society for Worldwide Interbank Financial Telecommunication (SWIFT) provides guidelines that the regulator can use to assess whether bank security is good enough to prevent certain known attacks. Damage-control investment, the success of which is hard to predict, can be regulated indirectly by liability rules. Liability rules, which are governed by the tort system, state the amount of damage each party is liable for. For example, software users may file lawsuits against software providers for security breaches, data leakage, and infringement of privacy, and if providers are proven to have caused harm, they will be held accountable for user damage.

Clearly, if the provider is not responsible for damage harming users, its investment incentives will be suboptimal. In the existing literature on bilateral care—where both the provider and the users can undertake one type of investment to reduce the expected damage—conventional wisdom suggests that strict liability with a defense of contributory negligence—under which the provider is fully liable only if the user is not negligent—yields the optimal investment (Brown, 1973). I, however, show that when the provider undertakes multiple types of investments, its investment incentives can still be suboptimal even when it has full liability. In particular, the provider underinvests in attack prevention and overinvests in damage control. The reason

is that the provider does not take into account the precautionary costs of the users, which gives it too much incentive to search for bugs. Moreover, because attack-prevention and bug-detection investments are substitutes, allowing providers to fix security problems later increases the likelihood of releasing a less secure software product in the first place. This result is akin to the practice of software versioning in the software industry, where providers first release versions of products that are prone to security issues and then fix these problems only at later stages (see Section 5.2). Interestingly, a partial liability rule (or more precisely, the provider bears a fine/reimbursement that is smaller than user damage level) with an optimal standard can restore the first-best outcome. And this result is consistent with the view taken by some security experts: Bruce Schneier, for instance, argued that

> *"100% of the liability should not fall on the shoulders of the software vendor, just as 100% should not fall on the attacker or the network owner. But today, 100% of the cost falls directly on the network owner, and that just has to stop."*[4]

The important implications of these results are that the regulator can implement similar standards of security as other, already regulated, industries such as automotive and aviation, and implement policies that help users reduce their costs of taking precautions. For instance, since not all users apply patches immediately after their introduction (home users may ignore security risk warnings, while enterprise users may not apply patches in a timely manner because of time constraints), policies that help synchronize patch release and adoption cycles can be useful (see Section 4.1). Furthermore, I show that increasing the number of expert users improves social welfare. On the other hand, it may exacerbate the under- and over-investment problems, which has important implications for user education in the software industry. The difference between private and social investment incentives arises from two sources of inefficiency. The first is that the provider does not pay fully for the damage, and the total amount of damage is decreasing in the number of expert users. The second source of inefficiency is that the provider ignores the precautionary costs of the users, and the total cost of precaution is increasing in the number of expert users. When the provider bears substantial liability for user damage, the second source of inefficiency dominates. These results suggest that if the objective of the government is to improve social welfare, it would be desirable to provide more support and training in the area of cybersecurity so that users become more competent in managing security threats. If its objective, however, is to alleviate inefficiencies in investments, then the government needs to be careful about increasing the number of expert users because the objectives of the social planner and the provider may become more divergent (see Section 4.2).

1.1 Paper contribution and literature

This paper extends bilateral care models to incorporate multiple types of investments on the part of the provider (namely, attack prevention and damage control), and to show that the joint use of a standard and partial liability can resolve investment inefficiencies. The result of a partial liability rule being optimal in this paper is in line with the results in the literature

[4]See Schneier (2007).

on asymmetric information, which studies how the presence of double moral hazard problem affects optimal warranty design. When the provider has private information about the quality of its product while this is unobservable to users, it can use warranties to signal good quality. However, if providers offer full warranties to users, the latter might not exercise reasonable care, which leads to a double moral hazard problem. Cooper and Ross (1985) and Belleflamme and Peitz (2010), for instance, show that under double moral hazard the optimal warranty calls for partial compensation for a defective product. I, however, show that a partial liability rule supports optimal care even when all investments are publicly observable (to the provider, the users and the courts). This means that the result of a partial rule being optimal does not hinge on moral hazard but rather on the presence of user precautionary costs. In that case, it is important for the regulator to recognize that solving the moral hazard problem is not sufficient to restore investment efficiency. Instead, the regulator should focus on the design of policies that share the burden of care between the provider and the users.

This paper further contributes to two strands of literature. First, it is related to recent works on the economics of security investment (for surveys, see Anderson, Clayton, and Moore, 2009; Anderson and Moore, 2009). Gordon and Loeb (2002) study the optimal protection of information, which varies with the information set's vulnerability.[5] Kunreuther and Heal (2003), August and Tunca (2006, 2011), Acemoglu et al. (2013), and Riordan (2014) study investment incentives in the presence of network externalities.[6] My model differs from these papers in that they consider each provider taking one action, whereas the provider in this paper can undertake both attack-prevention and damage-control investments. Varian (2004) examines full liability in a model in which efforts of multiple parties are needed to increase security. He finds that liability should entirely be assigned to the party who can best manage the risk. Contrary to his analysis, I also consider partial liability and the joint effect of partial liability and standards.

Second, although this paper relates to the economic and legal literature on tort laws,[7] it departs from traditional bilateral care models (see, e.g., Brown 1973; Shavell 1980; Landes and Posner 1985; Daughety and Reinganum 2013a) in the following way: whereas in the literature each party can take one type of care, in my model, the provider can invest in attack prevention and damage control, in addition to which the user can take precautionary action. Modeling in this way, I find that a partial liability rule yields the socially efficient outcome, which differs from what is found in Brown (1973).[8] The sources of the difference in results

[5]There are other security investment models in computer science (for a survey, see Böhme 2010), which, for instance, investigate questions about the appropriate amount of security budgets (i.e., how much to invest) and providers' security investment strategies (i.e., when and where to invest). They, however, do not tackle the investment problem from the legal and economic perspectives, meaning that the effects of security standards and liability policies on investment incentives (i.e., what measures should the regulator implement) have been largely ignored in this literature.

[6]As August and Tunca (2006) focus on the problem of patch management, they only consider damage-control investment. Security investments are strategic complements in Kunreuther and Heal (2003), strategic substitutes in Acemoglu et al. (2013), and are either the ones or the other in Riordan (2014) depending on whether the attacks are direct or indirect, but agents can only invest once in these models.

[7]See Shavell (2008) and Daughety and Reinganum (2013b) for excellent surveys of this literature.

[8]Since I do not consider usage in this model, Shavell (1980) and Landes and Posner (1985), who study

are explained in Section 3. Some literature also focuses on attack-prevention investment, such as Daughety and Reinganum (1995, 2006), or damage-control investment, such as Polinsky and Shavell (2010),[9] instead of both. Other papers such as Shavell (1984) and Kolstad et al. (1990) compare standards with liability rules. Shavell's analysis, however, is based on the inefficiencies associated with the provider's potential bankruptcy and the uncertainty of lawsuit by the users, while the inefficiencies studied by Kolstad et al. are due to the uncertainty over the legal standard to which the provider will be held liable. Here, though, the inefficiencies are caused by the presence of user precautionary costs and the possibility that the provider can allocate investments between attack-prevention and damage-control activities.

2 The model

Monopoly software provider. Consider a software provider that sells a software product which contains potential bugs that can be exploited by the hackers. To increase the security level of the software product s, the provider can invest $c(s)$ to reduce the probability of attacks to $p(s)$.[10] Such investment could take the form of improvement in infiltration detection or authentication technologies. In addition, the provider can invest $m(b)$ in damage control, which enables it to find bugs before the hacker does with probability b.[11] To focus on the main insights, I assume that upon discovering a bug, the provider discloses it to the users, who can then take precautionary actions such as patching (described below).[12] I assume away prices so that the problem is simplified to choosing a level of security that minimizes the sum of investment costs and expected damage (defined below). This assumption could be reasonable to the extent that the provider has to decide on the amount of investments after its sales, e.g., whether to install the upgrades and patches of different operating systems and software applications. Moreover, if the provider generates profit from channels other than selling the software product e.g. advertisement, then the objective is simply to minimize the costs.[13]

To facilitate the analysis, I make the following assumptions:

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

proportional-harm model (i.e., the effect of harm is linear on usage), and Daughety and Reinganum (2013a), who focus on cumulative-harm model (i.e., the effect of harm is non-linear on usage), are not the primary point of comparison with this model.

[9] Polinsky and Shavell analyze information acquisition about product risks when product quality is uncertain. Therefore, their problem concerns damage-control, rather than attack-prevention, investment.

[10] I assume away strategic attacks, which 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.

[11] Whether the provider chooses s and b sequentially or simultaneously does not affect the results, but in practice attack prevention and damage control usually happen sequentially. In any case, the novel feature of this model is to have multiple types of investments on the part of the provider.

[12] Section 5.3 examines the implications of bug disclosure and bug bounty programs.

[13] This model focuses on the problems associated with two types of investments and means to address these problems, and so abstracts away from prices. For a related discussion of the effect of prices but without different types of investments, see, e.g., Daughety and Reinganum (2013a).

Under Assumption 1, investment costs $c(s)$ and $m(b)$ are increasing and convex in s and b respectively;[14] the probability of attack $p(s)$ is decreasing and convex in s; and $c(s)$ and $p(s)$ are thrice differentiable in s.[15]

Heterogeneous software users. There is a unit mass of software users, who can be of two types: a proportion α of them are "experts" and have precaution cost γ drawn from a distribution $F(\gamma) \sim [0, +\infty)$, while the others are "laymen" with $\gamma = \infty$.[16] Both α and $F(\gamma)$ are common knowledge. Experts have better security knowledge and awareness and can take precautions such as monitoring the system for attacks and patching the system if the provider finds and discloses a bug, while laymen without such knowledge will never take precautions.[17] The analysis applies more generally to both enterprise as well as home users. Using the example in the Introduction, Dropbox was attacked recently, which caused damage to both enterprise and home users. Both types of users can take precautionary action against the attacks, although their precautionary behavior may vary. For example, home users can adopt multi-factor authentication (e.g., user login requires not only entering a password but also a code that is sent to the users' mobile phone), whereas enterprise users can adopt user-side encryption (e.g., they can encrypt the more sensitive files before storing them in the cloud).

I assume that all investments (including the provider's investments and user precautionary actions) are publicly observable. That is, I chose not to model the moral hazard problem because this setup enables me to highlight the source of investment inefficiency resulting from both parties investing and one of them investing in multiple activities rather than the moral hazard problem itself (the latter of which has been emphasized in the literature on warranty design), which suggests that policies that merely get rid of the moral hazard problem are not enough to restore efficiency.

Damage. When an attack occurs, the provider incurs a damage of $\overline{\eta}$ if the hacker discovers the bug before the provider does, and $\underline{\eta}$ if the provider identifies the bug first. This could be financial losses and reputational harm caused by information stolen from the provider. Such loss is smaller if the provider finds the bug first as it can then try to fix the problem. The provider, however, may face substantial loss if the hacker exploits a bug that has not been previously identified—a phenomenon known as "zero-day attacks". That is, $\overline{\eta} > \underline{\eta}$. As for the users, they incur a damage of $\overline{\mu}$ 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. User precautionary action can limit the extent of damage resulting from an attack, so $\overline{\mu} > \underline{\mu}$.

Policy instruments. There are two policy instruments available. First, the regulator can

[14] If there are externalities between the two cost functions, meaning investing more in attack prevention will make finding bugs easier, then there will be more investment in attack prevention under both optimal and equilibrium regimes because of cost reductions in damage control. However, it will not change the qualitative result that partial liability rule supports optimal investment, provided the software provider does not take into account user precautionary cost when choosing its investments.

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

[16] As discussed in Section 4.2, α is included in the model because it yields interesting implications for user education.

[17] I assume that users take precaution after the provider has discovered and disclosed the bug. One could alternatively think of users 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 provider.

set a minimal security standard, s^R, for attack-prevention investment. It is easier for the regulator to know the optimal level of s than that of b because standards are usually set to prevent known vulnerabilities. For example, the regulator can set up some security tests that a software product needs to pass in order to be released on the market. I, however, assume that the regulator cannot regulate damage-control investment directly: the probability of a software provider finding a bug earlier than the hackers, b, is difficult to monitor and predict, as it largely depends on the constantly evolving technologies of both the hackers and the defenders. And at any point in time, there can be new zero-days.[18]

Second, the regulator can impose a liability rule $\lambda \in [0, 1]$, where λ denotes the part of user damage for which the provider is liable. In reality, liability can be imposed as a fine paid by the provider to the regulator, in which case the fine does not affect user precautionary behavior. Alternatively, liability can be imposed as a reimbursement to the users, in which case the refund does affect user precautionary behavior. For example, fines are common in the IT industry. Regulatory bodies such as the British Information Commissioner's Office can issue fines to providers that breach the UK Data Protection Act. Companies such as Sony and eBay have historically been fined for a breach of the Act. Another example is AT&T's data breaches in 2013 and 2014 occurring at three of its international call centers. This led to a $25 million fine, which is the largest penalty ever imposed by the Federal Communications Commission on a company for data security and privacy violations.[19] Reimbursements, however, are more common in other industries. For instance, many banks take full liability for unauthorized online transactions, meaning users are reimbursed for all financial losses originating from identity theft. US retailer Target provides another example.[20] It will reimburse the victims of its data breach, which occurred in 2013 and resulted in the theft of at least 40 million credit card numbers.[21] More generally, fines are more applicable to cases where it is difficult for users to file lawsuits (e.g., because of the triviality of the security breach or the lack of financial resources to go against big firms), so that the provider might not even be able to identify the victims of the attack to offer a refund. We begin the analysis by using the interpretation of liability as a fine, and in Section 5.1 we discuss the interpretation as a reimbursement. Yet, the main result of Proposition 2 (below) that the partial liability rule yields the socially efficient outcome holds under both interpretations. Moreover, the main result remains valid in an alternative model with a continuum of users instead of two types considered here (see Appendix A for the case with fine and Appendix H for the case with reimbursement).

[18]Symantec (2016) finds that in 2015 a new zero-day was discovered each week on average.

[19]See "AT&T pays record $25 m fine over customer data thefts," *BBC News*, April 9, 2015, available at `http://www.bbc.com/news/technology-32232604`.

[20]In this example, Target is the software provider and shoppers at Target are the software users. Target manages a technology that stores its customers' credit card information. Hence, in the event of an attack against Target (owing to its website's insecurity), Target but not the software companies or the credit card companies has to make the reimbursements. Note that, to keep the analysis tractable, addressing more complex cases involving more than two parties (the provider and the users) is beyond the scope of this analysis. An example is some credit card frauds that affect banks as well as credit card companies, merchants, and cardholders.

[21]See "Target to pay $10 m to settle lawsuit over data breach," *BBC News*, March 19, 2015, available at `http://www.bbc.com/news/technology-31963612`.

There are three possible liability regimes:[22]

- Full liability, under which the provider is liable for all damage incurred by the users, i.e., $\lambda = 1$;
- Partial liability, under which the provider is partially liable for user damage, i.e., $\lambda \in (0, 1)$; and
- No liability, under which the provider is not liable for any user damage, i.e., $\lambda = 0$.

Thus, the total loss for the provider is $\eta + \lambda\mu$, where $\eta \in [\overline{\eta}, \underline{\eta}]$ and $\mu \in [\overline{\mu}, \underline{\mu}]$. Notice that full liability here is defined for "net" damage to user, i.e., μ. One can alternatively define it for "total" damage, which includes also user precaution cost γ. I model the liability regime in this way because, in practice, providers are typically liable for financial damages to the users caused by, for example, identity theft losses related to a data breach. Liability sometimes also covers litigation costs but rarely investment costs in precaution.[23] One difficulty lies in verifying the amount of time and effort users spent on managing, maintaining and patching a system.

In sum, we consider the following game: first, the social planner chooses a security policy (a standard and/or a liability rule). Then, the software provider makes its investment decisions in attack prevention and damage control. Finally, if the provider discovers and discloses a bug, users decide whether or not to take precautionary actions.

3 Investment incentives

The user's problem

Proceeding backward, we start with the user's problem. When there is no disclosure about bugs to the users, they cannot do anything. When the provider finds and discloses a bug, the expected damage for a user who does not take precaution is $p(s)\overline{\mu}$, and that for a user who takes precaution is $p(s)\underline{\mu} + \gamma$. Therefore, the user will take precaution if

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

[22]The legal literature uses other terminologies, for instance, they call the situation wherein a provider must fully compensate a user for harm caused "strict liability" instead of "full liability"; the situation wherein a provider is liable only if the user is not negligent "strict liability with a defense of contributory negligence"; and the situation wherein a negligent provider is only partially liable if the user is also negligent "comparative negligence"(See Brown (1973) pp. 328-331 for a clear description of different liability rules). Much of the legal literature, however, focuses on the first two rules, but rarely discusses the role of partial liability.

I adopt slightly different terminologies to disentangle the effect of the two instruments—standards and liability rules—because in practice they are usually implemented by separate regulatory agencies. For example, rather than one court or agency making a centralized decision altogether, we have the tort system governing the circumstances under which a party is liable for damage caused, and other regulatory agencies setting standards.

[23]Incorporating litigations in the model would not change my qualitative results since in the alternative model the expected damage faced by users will change from μ to the probability of losing the litigation times μ, whereas the provider's expected liability will become its probability of losing the litigation times λ times μ; but all that matters is the magnitude.

The provider's problem

In the investment stage, the provider chooses s and b to minimize the sum of expected damage and investment costs, which is denoted by L^f. That is, the provider's problem is

$$\min_{b,s} L^f = \underbrace{(1-b)p(s)(\overline{\eta}+\lambda\overline{\mu})}_{hacker\ finds\ the\ bug\ first} + \underbrace{b\left\{\int_0^{p(s)(\overline{\mu}-\underline{\mu})} p(s)[\underline{\eta}+\lambda(\alpha\underline{\mu}+(1-\alpha)\overline{\mu})]dF(\gamma)\right.}_{provider\ finds\ the\ bug\ first\ and\ experts\ take\ precaution} + \underbrace{\left.\int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} p(s)(\underline{\eta}+\lambda\overline{\mu})dF(\gamma)\right\}}_{no\ users\ take\ precaution} + \underbrace{m(b)+c(s)}_{investment\ costs}. \quad (2)$$

The first term in Equation (2) is the expected cost of the provider when the hacker discovers the bug first, in which case both the provider and the users suffer a large damage, $\overline{\eta}$ and $\overline{\mu}$, respectively. When the provider finds the bug before the hacker does and discloses the bug, it suffers a small damage $\underline{\eta}$, whereas there are two cases for the users. In one, when expert users have low precautionary costs such that Equation 1 is satisfied, they will take precautions, hence reducing their damage to $\underline{\mu}$. As for laymen, they never take precautions due to the high costs, and they will suffer a large damage $\overline{\mu}$. This is captured by the second term. In the other, the case where expert users have high precautionary costs, no users will take any precautions and all of them will incur large damage when an attack occurs, and this is captured by the third term. The last two terms represent the costs of attack-prevention and damage-control investments. Let $b^m(s)$ denote the provider's optimal damage-control investment strategy given attack-prevention investment s, and let s^* and $b^* \equiv b^m(s^*)$ denote the solutions of Equation (2).

The social planner's problem

Turning to the social planner's problem, after discovering a bug, it is always better for the social planner to disclose the bug, as this induces some users to take precautionary actions and reduces the total damage when an attack occurs. As above, when a bug is disclosed, expert users take precautions when the cost of doing so is small and do not take precautions when the cost is large. The social planner's problem then is to choose s and b to minimize the expected costs of the society, which is denoted by L^{SP}:

$$\min_{b,s} L^{SP} = (1-b)p(s)(\overline{\eta}+\overline{\mu}) + b\left\{\int_0^{p(s)(\overline{\mu}-\underline{\mu})} [p(s)(\underline{\eta}+\alpha\underline{\mu}+(1-\alpha)\overline{\mu})+\alpha\gamma]dF(\gamma) + \int_{p(s)(\overline{\mu}-\underline{\mu})}^{\infty} p(s)(\underline{\eta}+\overline{\mu})dF(\gamma)\right\} + m(b)+c(s). \quad (3)$$

There are two differences with the provider's problem. First, in the case where a bug is disclosed, the social planner takes into account the impact of user precautionary costs ($\alpha\gamma$)

on investments, whereas the provider does not. Second, the social planner internalizes all the society costs, so there is no liability issue. Let $b^{SP}(s)$ denote the social planner's optimal damage-control investment strategy given attack-prevention investment s, and let s^o and $b^o \equiv b^{SP}(s^o)$ denote the solutions of Equation (3).

Moreover, we can rewrite the objective function of the social planner as follows:

$$L^{SP} = L^f|_{\lambda=1} + \underbrace{b\alpha \int_0^{p(s)(\overline{\mu}-\underline{\mu})} \gamma dF(\gamma)}_{users' \ costs}. \tag{4}$$

It should then be clear that the difference between L^f and L^{SP} is due to the fact that the provider minimizes its own private costs, while the social planner minimizes the sum of the provider's and the users' costs.

Investment Inefficiencies

Let us now discuss how the provider's investment incentives may diverge from those of the social planner and the optimal policy that can address this problem. First consider what happens when the provider bears all the liabilities.

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 provider has less incentive to find bugs given a high security level for attack prevention, meaning that attack-prevention and damage-control investments are substitutes. The next lemma characterizes damage-control investments under full liability and an optimal standard. Specifically,

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 provider will overinvest in damage control,* $b^m(s^o) > b^{SP}(s^o)$*.*

Proof. See Appendix C. □

One might expect that under full liability and an optimal standard the provider will invest optimally, yet this is not the case when users also bear some precautionary costs. The intuition runs as follows. If a bug is not found by the provider, the provider and society suffer the same magnitude of heavy losses because hackers can exploit the bug fully before it is patched by the provider and since the bug is not identified, users cannot take any precautionary actions to reduce damage. If a bug is discovered by the provider, it suffers less damage than the social planner because once the problem is disclosed users can employ various defenses against the attacks. However, since the provider decides on the amount of investment to minimize its own private costs, it will ignore the precautionary costs on the part of the users, whereas the social planner minimizes the sum of these costs. We can see this clearly from Equation (4). Because the provider has more to gain in finding bugs, it will overinvest with respect to the socially efficient level.

We can further show that the full liability regime is suboptimal:

Proposition 1. *(Full liability with fines). Under full liability* ($\lambda = 1$)*, under which the provider must pay a fine to the regulator for all the damage caused, the provider underinvests in attack prevention,* $s^* < s^o$*, and overinvests in damage control,* $b^* > b^o$*.*

Proof. See Appendix D. □

As we can see from Proposition 1, full liability alone does not achieve the first-best solution. The reason is that, as shown in Lemma 2, the provider has more to gain in finding a bug than the social planner, which results in too much investment in damage control. This, in turn, results in too little investment in attack prevention (because s and b are substitutes, as shown in Lemma 1). Furthermore, in Appendix F, I show that if liability regime is the only instrument of regulatory policies, no liability regime (whether full, partial or zero) suffices to provide the right incentives for two investments. Instead, a standard regulation with an appropriately chosen partial liability rule can implement the first best.

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* $s^R = s^o$ *and an optimal partial liability rule* $\lambda^o \in (0, 1)$*.*

Proof. See Appendix E. □

When security standards are set at the socially optimal level, it is inefficient to implement full liability because the provider will overinvest in damage control; it is also inefficient to set the provider's liability to zero because it will then underinvest in damage control. As a consequence, the optimal liability rule is a partial one:

$$\lambda^o = 1 - \frac{\int_0^{p(s^o)(\bar{\mu}-\underline{\mu})} \gamma dF(\gamma)}{p(s^o)(\bar{\mu}-\underline{\mu})F(p(s^o)(\bar{\mu}-\underline{\mu}))} \in (0,1),$$

which depends on the damage prevented by taking precaution $p(s^o)(\bar{\mu}-\underline{\mu})$ and the distribution of precautionary costs $F(\gamma)$.[24]

This result is related to Brown's (1973) work on bilateral care model, wherein he finds that "strict liability with a defense of contributory negligence" (i.e., the provider is fully liable unless the user is negligent) supports the socially efficient outcome. On the contrary, I find that full liability, with and without a standard, does not achieve efficiency; instead partial liability and an optimal standard do, both for the cases with fine (this section) and with reimbursement (Section 5.1). The main difference between Brown's model and this one is that in his model both parties (the provider and the user) choose one type of investment, whereas in this model one party (the provider) chooses two types of investments. In his model, under strict liability with a defense of contributory negligence, both parties will exert the right level of care: as the provider is liable for all losses sustained by users, it is induced to choose the correct level of care. Moreover, since users know that they have to bear all the losses if they are negligent (i.e., when they fail to meet the due care standard), they will exercise the correct level of care. With multiple types of investments, however, the provider can reallocate investment from one activity to another and in that case, investment inefficiencies can result even under full liability.

[24] In the simplest case where γ is uniformly distributed, we have $\lambda^o = 1/2$, which shows that the optimal liability can be significantly different from full liability.

4 Policy implications

Let us now discuss the implications of these results for regulatory policies and education in cybersecurity.

4.1 Standards and partial liability

Proposition 2 shows that security can be improved with the joint use of an optimal standard and a partial liability rule. For standards, they can either be implemented by the legal system in the form of negligence rules, under which the party who does not comply with the due care standard chosen by courts will be penalized, or by a separate regulatory agency. Such agency can establish minimum standards for IT security (such as a mandatory compliance framework in encryption and security breach notification) as other already regulated industries like automotive and aviation, where new models of car and aircraft must pass some safety tests conducted by international or national regulatory bodies before they are allowed on the road or in the air.

As for liability rules, they can be implemented by the tort law system. The result that partial liability and a standard can achieve optimality implies that the regulator should not impose a one hundred percent liability on the software provider, as also suggested by some security experts, e.g., Bruce Schneier. Instead, an effective policy is to ask both the software provider and its users to share the costs of security.

Furthermore, since investment inefficiencies result from the presence of user precautionary costs, it immediately follows from Equation (4) that

Corollary 1. *(User precautionary costs). When there are no precautionary costs ($\gamma = 0$), full liability ($\lambda = 1$) implements the efficient outcome, i.e., $s^* = s^o$ and $b^* = b^o$.*

Hence, it would be desirable from society's point of view to reduce user precautionary costs if the regulator were to impose 100% liability on the provider. This has implications for both home and enterprise users. As mentioned in the introduction, poor security habits of both types of users (e.g., being slow to patch) bear low-hanging fruit for hackers.

At an individual level, despite the fact that home users dislike or feel concerned about security problems, many of them do not take appropriate care, e.g., they do not patch their machines and use simple passwords. One reason for this may be that users have other competing demands on their time, and paying attention to security issues appears to be low on their priority list. Therefore, it can be useful to simplify user precautionary measures such as releasing automatic patching. Moreover, because of the hassle of remembering strong and multiple passwords, many users use easy-to-remember passwords and reuse the same credentials across websites. Thus, another way of reducing precautionary costs is to promote the development and adoption of password managers, which can generate and store unique passwords, thereby saving on user hassle costs.

At an enterprise level, installing patches could be time- and resource-consuming, especially for large companies. Indeed, the plethora of security updates can often overwhelm software engineers, as they have to keep track of all relevant bugs and patches, as well as match the

version of all these updates to those of the software used by their company. Once a problem is identified, they need to figure out which updates get priority and look for solutions.[25] As a consequence, few companies can apply updates in a timely manner, which may induce major security problems. This suggests that a desirable policy should try to eliminate the delay in applying solutions to security problems. First, it can be useful to encourage providers to release automatic security updates.[26] Second, third parties can be introduced to help enterprises to find, select and deploy the solutions that are relevant to their systems. For example, using the cloud computing technology, enterprise users could outsource security activities to external cloud providers. The establishment of vulnerability management companies could also be useful in helping enterprise users to adhere to compliance and security standards.

4.2 User education

Since it is the proportion of experts able to take precautions that drives the difference between the provider's investments and those of the social planner, it would be interesting to examine how an increase in α would change investment incentives:

Corollary 2. *(User education). When λ is large, increasing the proportion of experts, α, exacerbates underinvestment in attack prevention and overinvestment in damage control.*

Proof. See Appendix G. □

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

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

Investment incentives are distorted by two forces: first, the provider does not pay fully for the damage; second, the provider ignores user precautionary costs. If the provider 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 provider's liability reduces the first type of distortion, whereas a decrease in the proportion of experts reduces the second one. Taken together, the objectives of the social planner and the provider become more aligned, and thus a decrease in α reduces the extent that the provider is investing suboptimally.[27] Thus, if the goal is to alleviate investment inefficiency, the government needs to be careful about increasing the number of experts.

[25] Practitioners have commonly considered patch management as a time- and resource-consuming activity. See "Automating Patch Management," *Symantec*, February 8, 2005, available at `http://symc.ly/1KHNGrS`. On average, enterprise users spend 600 hours per week on malware containment processes. See "Four in five malware alerts are a 'waste of time'," *ZDNet*, January 19, 2015, available at `http://zd.net/1MWu7Oj`.

[26] L^f and $L^S P$ are increasing in γ because more users taking precaution lowers the expected damage, so the provider has incentives to lower γ. However, these incentives are still suboptimal when $\lambda = 1$.

[27] If λ is small, the provider, knowing that they are only liable for a small part of the damage, may underinvest in both attack prevention and damage control. The effect of α is then difficult to generalize because it depends on not only λ, but also s^*, s^o, b^* and b^o.

Nevertheless, this does not mean that offering cybersecurity training is undesirable from society's point of view. More specifically, Equation (3) shows that an increase in the proportion of experts leads to a decrease in society's loss. This suggests that if the primary objective of the government is to improve social welfare, policymakers can provide support and training in the area of cybersecurity so that users become more competent in managing security threats. For example, many security breaches involve hackers trying to compromise user accounts and users may be unaware of such attacks. Even if they are, they may lack the necessary skills to solve the security problem. Therefore, increasing training that aims at enhancing the skills of these users can be useful. The government, however, should keep in mind that increasing the number of experts involves a trade off between loss reduction and investment inefficiencies.

5 Further discussion

This section discusses alternative interpretations of this model and, in particular, how the underinvestment and overinvestment results can be used to explain real-world security issues in IT and in other industries where providers undertake two types of investments. It also considers how to address these issues by implementing alternative policies, such as reimbursing users instead of imposing fines.

5.1 Reimbursement

So far, we have interpreted liability as a fine, which does not affect user precautionary behaviors. Now suppose that the provider, instead of paying a fine to courts for a proportion λ of user damage, is required to reimburse users an amount specified by one of the liability regimes. Let ρ denote the refund that returns to the pockets of users, which can be equal to or smaller than user damage level. I show that

Proposition 3. *(Full liability with reimbursements). Under full refund ($\rho = 1$), under which the provider must reimburse fully to a user for damage caused, the provider overinvests in attack prevention, $s^* > s^o$, and underinvests in damage control, $b^* < b^o$.*

Proof. See Appendix H. □

Under full reimbursement, if a bug is not found by the provider, the provider and society suffer the same loss. However, if a bug is found and disclosed, users, knowing that they will be fully reimbursed anyway, have no incentive to take precaution in equilibrium; under social optimum, though, some users (the experts in particular) will take precaution when the benefit of an increase in precautionary action outweighs the cost. The benefit of investing in b for the provider therefore is less than that of the social planner. Consequently, the provider underinvests in damage control. And since attack-prevention and damage-control investments are substitutes, the provider overinvests in attack prevention.

The banking sector is a case in point. Financial institutions invest a large amount of money in developing new technologies that defend their users against password theft, but much less

in damage reduction because tracing suspicious money transfers from one bank account to another is relatively easier than preventing password-stealing attacks in the first place.

It is also straightforward to show that Proposition 2 remains valid in the case with reimbursement. Although different instruments (a fine or a reimbursement) yields different investment incentives, in both cases a partial liability rule results in the socially efficient outcome. This suggests that when both the provider and the users can invest in security and the provider can undertake two types of investments, policymakers can think about conferring some liability to users instead of adopting either full or zero liability rules. Likewise, they should consider whether a fine or a reimbursement is the more appropriate regulatory instrument.

5.2 Software versioning

The result of underinvestment in attack prevention and overinvestment in damage control is consistent with the current development in the software industry, where software providers often "experiment" the alpha versions of their products (e.g., software, mobile applications, and smart-home appliances) in public and release improved beta versions at a later date. Thus, alpha versions of many software products are susceptible to security risks. The result of underinvestment in attack prevention in this model captures the essence of this situation. The new insight is that the possibility of sequential investments, which allows the provider to fix security problems later, provides an alternative if partial explanation for why the provider may release software products prematurely.[28]

5.3 Bug bounty programs

So far, I have assumed that the software provider discloses a bug if it finds the bug before the hacker does. This assumption does not affect the main insights: even if the provider has the option of not disclosing the bug, it will never do so; the provider always weakly prefers to disclose so as to induce the expert users to take precautions.[29]

This model has also an interesting implication for bug bounties, which are financial rewards offered by software providers to users for reporting bugs. More specifically, bug bounty programs can have negative effects on investment incentives because of the presence of multiple investments: they give more incentives for users to invest, which may lead the provider to overinvest in bug fixing and underinvest in attack prevention even more (i.e., a very bad alpha version may be released). This is in contrast to Choi et al. (2010a), who find that bug

[28]My results are also related to the literature on "vaporware," the practice of which refers to situations where providers announce new products well in advance of their actual release on the market. If we view the announced product as a product characteristic (e.g., a security feature), then vaporware could mean delivering a lower-quality product compared to what was claimed by the provider, which is akin to the result of underinvestment. Interestingly, I show that the practice of vaporware can be a profit maximizing strategy for the provider as long as it undertakes multiple types of investments, whereas the previous literature finds that providers engage in vaporware only to prevent entry (see Bayus et al., 2001; Haan, 2003) or when reputational concern is not so important (see, Choi et al., 2010b).

[29]The provider strictly prefers to disclose it when expert users' precautionary costs are low (i.e., when Equation (1) holds) and is indifferent between disclosing or not when precautionary costs are high.

bounty programs can be welfare-improving due to an increase in disclosure by the provider.[30] Although such disclosure-increasing effect is absent in this paper, the potential adverse effects of bug bounty programs emphasized here should not be ignored. In particular, if the investment inefficiencies caused by the presence of user precautionary costs are significant, then a more effective policy to improve efficiency is to implement standards and a partial liability rule instead of bug bounty programs.

5.4 More general applications

This analysis also provides insight into other industries in which sequential investments (e.g., a first investment in pre-sale product design and a second investment in post-sale remedial measures) are made by firms that produce a product with some safety features. For example, in automobile production, the pre-sale investment could lead to the development of a new technology in cars that is subject to some potential safety defects. After sale, the firm can invest in remedying these safety problems. It is, for instance, common to observe product recalls because of problems in engines or braking in the car industry. Workplace and factory design provide other examples. Employers first have to invest in safety technologies to reduce work hazards and prevent injuries, and then in medical care if any injuries occur and equipment replacement. Likewise, factory owners first have to determine a level of precaution to minimize health and environmental risk (e.g., pollution and radiation), and then a certain inspection frequency of the facility. With these re-interpretations, this model might be useful to analyze the investment incentives of the various stakeholders and, more specifically, whether there are incorrect incentives on the part of the injurer and the potentially injured, and if so, how to correct them.[31]

6 Conclusion

More and more devices, such as mobile phones, home appliances, health devices, cars, and even some infrastructures (e.g., traffic lights), become Internet connected, but we continue to discover security failures, including malware, poor encryption and backdoors that allow unauthorized access. This paper suggests that to increase security, the key is not so much to hold the provider of these devices solely liable for the loss as to balance the investment incentives between the provider and the users.

In practice, there are few policies regulating the software industry compared to financial services and transportation. Establishing national or international regulatory body to implement security standards and update existing regulations to ensure that only products with adequate defenses against attacks can be released on the market could represent a useful start.

[30] Another difference between this paper and Choi et al.'s is that they take security investments as given and do not discuss optimal investments, which is the focus of this paper.

[31] However, note that these examples may not fit my model along all aspects. For example, software versioning is more common in IT, whereas there is generally a stricter compliance framework for firms in the other examples like automobiles and workplace design.

For example, Finland passed a new legislation, the "Information Society Code", at the beginning of 2015, which enforces security standards on a wide range of platforms such as Apple, Facebook and Twitter.

In future work, it would be interesting to relax the single-provider assumption and study competition between software providers. The possibility of interdependencies between software products may lead to interesting dynamics between providers, and investment incentives may be different depending on whether providers' investments are substitutes. Alternatively, one could study contagion issues in a network of multiple providers.[32]

Appendices

A Continuum of users

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

If the provider does not find a bug before the hacker does, then its expected cost is $p(s)(\overline{\eta} + \lambda\overline{\mu})$, whereas if it finds and discloses the bug, then its expected cost is $p(s)[\underline{\eta} + \lambda(F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu})]$. Hence, the provider chooses s and b to minimize

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

As for the social planner, if a bug is not found, then the expected cost is $p(s)(\overline{\eta} + \overline{\mu})$, whereas if the bug is found and disclosed, then the expected cost is $p(s)[\underline{\eta} + F(\gamma(s))\underline{\mu} + (1 - F(\gamma(s)))\overline{\mu}] + \int_0^{\gamma(s)} \gamma dF(\gamma)$. The social planner therefore solves

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

It is easy to see that since $\int_0^{\gamma(s)} \gamma dF(\gamma) > 0$, $L^{SP} > L^f$ for any λ. Thus, the main results of underinvestment in attack prevention and overinvestment in damage control carry through.

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

B Proof of Lemma 1

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

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

and

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

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

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

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

and

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

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, and the second inequality from 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 $H^{SP}(s)$ and $H^f(s)$ 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 $H^{SP}(b^{SP}(s)) > H^f(b^m(s))$.[33] Finally, using Lemma 1, if $s^* < s^o$, then $b^* > b^o$.

E Proof of Proposition 2

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

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

Comparing with Equation (B.1), we can see that $b^m(s^o) < b^{SP}(s^o)$. Since $b^m(s)$ is increasing in λ, there exists an optimal liability rule $\lambda^o \in (0,1)$ such that $b^m(s^o) = b^{SP}(s^o)$. More specifically, by equating $\partial L^{SP}/\partial b|_{s=s^o}$ and $\partial L^f/\partial b|_{s=s^o}$, we obtain the optimal liability rule:

$$\lambda^o = 1 - \frac{\int_0^{p(s^o)(\bar{\mu}-\underline{\mu})} \gamma dF(\gamma)}{p(s^o)(\bar{\mu}-\underline{\mu})F(p(s^o)(\bar{\mu}-\underline{\mu}))}.$$

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 $\partial L^f/\partial b = \partial L^{SP}/\partial b$ and $\partial L^f/\partial s = \partial L^{SP}/\partial s$. However, we can easily verify that these two conditions cannot be satisfied at the same time for the same λ.

G Proof of Corollary 2

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 overinvestment in damage control.

[33] 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.

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 underinvestment in attack prevention is more severe as α increases. This proof remains valid as long as λ is large enough.

H Proof of Proposition 3

First, with reimbursement, the problem for the social planner is the same as that in the case with fine. The social planner will ask the experts to take precaution if the cost of an increase in precautionary action is less than the benefit of reducing damage, i.e., when Equation (1) is satisfied.

In the market equilibrium, when users are reimbursed fully for all damage, their incentives to take precaution are weakened. More specifically, a user now takes precautionary action if

$$\gamma < (1-\rho)p(s)(\bar{\mu}-\underline{\mu}).$$

Thus, the provider chooses s and b to minimize

$$\min_{b,s} L_r^f = (1-b)p(s)(\overline{\eta}+\rho\overline{\mu}) + b\left\{\int_0^{(1-\rho)p(s)(\overline{\mu}-\underline{\mu})} p(s)[\underline{\eta}+\rho(\alpha\underline{\mu}+(1-\alpha)\overline{\mu})]dF(\gamma)+\int_{(1-\rho)p(s)(\overline{\mu}-\underline{\mu})}^{\infty} p(s)(\underline{\eta}+\rho\overline{\mu})dF(\gamma)\right\} + m(b)+c(s), \quad \text{(H.1)}$$

where subscript r denotes the case of reimbursement. The difference between Equation (2) in the main text (the case with fine) and the equation above (the case with reimbursement) lies in the boundaries of the integrals.

The first-order condition with respect to b (when $\rho = 1$) for the provider is

$$m'(b) = p(s)(\overline{\eta}+\overline{\mu}) - \int_0^{\infty} p(s)(\underline{\eta}+\overline{\mu})dF(\gamma). \quad \text{(H.2)}$$

In comparison with the first-order condition with respect to b for the social planner (see Equation (B.1)), it is clear that for a given s, the marginal benefit of investing in b for the provider is always lower than that of the social planner. Therefore, the provider underinvests in damage control.

As for the incentive to invest in attack prevention, the first order condition with respect to s for the provider is

$$-\frac{c'(s)}{p'(s)} = (1-b)(\overline{\eta}+\overline{\mu}) + b(\underline{\eta}+\overline{\mu}). \tag{H.3}$$

Comparing it with the first-order condition with respect to s for the social planner (see Equation (D.1)), it is easy to see that for a given b, the right hand side of the Equation (H.3) for the provider is always higher than that of the social planner. Therefore, the provider overinvests in attack prevention.

We can easily see that these results carry over to the case with a continuum of users with different precaution costs (which is equivalent to setting $\gamma(s) \equiv (1-\rho)p(s)(\overline{\mu}-\underline{\mu})$ in Appendix A).

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

[5] Terrence August and Tunay Tunca. Who Should Be Responsible for Software Security? A Comparative Analysis of Liability Policies in Network Environments. *Management Science*, 57(5):934–959, 2011.

[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] Paul Belleflamme and Martin Peitz. Marketing tools for experience goods. In *Industrial Organization: Markets and Strategies*, chapter 13, pages 309–330. Cambridge University Press, 2010.

[8] Rainer Böhme. Security Metrics and Security Investment Models. In Isao Echizen, Noboru Kunihiro, and Ryoichi Sasaki, editors, *Advances in Information and Computer Security, Lecture Notes in Computer Science*, volume 6434, pages 10–24. Springer Berlin Heidelberg, 2010.

[9] John Brown. Toward an Economic Theory of Liability. *Journal of Legal Studies*, 2(2):323–349, 1973.

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

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

[12] Russell Cooper and Thomas Ross. Product Warranties and Double Moral Hazard. *RAND Journal of Economics*, 16(1):103–113, 1985.

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

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

[15] Andrew Daughety and Jennifer Reinganum. Cumulative Harm, Products Liability, and Bilateral Care. *American Law and Economics Review*, 15(2):409–442, 2013a.

[16] 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., 2013b.

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

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

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

[20] 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.

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

[22] William Landes and Richard Posner. A Positive Economic Analysis of Products Liability. *Journal of Legal Studies*, 14(3):535–567, 1985.

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

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

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

[26] Bruce Schneier. Information Security and Externalities, 2007. Available at http://bit.ly/1MkAZo3 (accessed 20 October, 2016).

[27] Steven Shavell. Strict Liability versus Negligence. *Journal of Legal Studies*, 9(1):1–25, 1980.

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

[29] Steven Shavell. Liability for Accidents. The New Palgrave Dictionary of Economics, 2008. Available at http://www.law.harvard.edu/faculty/shavell/pdf/124_liability.pdf (accessed 26 October, 2014).

[30] Symantec. Internet Security Threat Report, 2016. Available at https://www.symantec.com/content/dam/symantec/docs/reports/istr-21-2016-en.pdf (accessed 6 September, 2016).

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