

Network Security: Vulnerabilities and Disclosure Policy[#]

by

Jay Pil Choi*, Chaim Fershtman**, and Neil Gandal***

October 22, 2007

Abstract

Software security is a major concern for vendors, consumers, and regulators since attackers that exploit vulnerabilities can cause substantial damages. When vulnerabilities are discovered after the software has been sold to consumers, the firms face a dilemma. A policy of disclosing vulnerabilities and issuing updates protects only consumers who install updates, while the disclosure itself facilitates reverse engineering of the vulnerability by hackers. The paper considers a firm that sells software which is subject to potential security breaches. Prices, market shares, and profits depend on the disclosure policy of the firm. The paper derives the conditions under which a firm would disclose vulnerabilities. It examines the effect of a regulatory policy that requires mandatory disclosure of vulnerabilities and shows that a 'Mandatory Disclosure' regulatory policy is not necessarily welfare improving. The paper then discusses the incentives to invest in product security. An ex-ante reduction in the number of vulnerabilities typically leads to higher prices, greater profits, and higher welfare, but may also induce a (welfare-improving) regime shift from a disclosure to non-disclosure policy. Ex-post investment may induce a (welfare-improving) regime shift in the opposite direction: from non-disclosure to disclosure.

JEL Classification: L100, L630.

Keywords: Internet security, software vulnerabilities, disclosure policy.

#We are grateful to Sagit Bar-Gill for excellent research assistance and thank Jacques Lawarree, Shlomit Wagman, and participants from the DIMACS 2007 conference and the UBC 2007 Summer Conference on Industrial Organization for their helpful comments. A research grant from Microsoft is gratefully acknowledged. Any opinions expressed are those of the authors.

* Michigan State University, e-mail: choijay@msu.edu

** Tel Aviv University, Erasmus University, and CEPR, e-mail: fersht@post.tau.ac.il

*** Tel Aviv University and CEPR, e-mail: gandal@post.tau.ac.il

1. Introduction

The Internet provides many benefits, but at the same time it also poses serious security problems. According to a study conducted by America Online and the National Cyber Security Alliance (2004), 80 percent of the computers in the US are infected with spyware and almost 20 percent of the machines have viruses. Some of these viruses have been very costly. According to the *Economist*, the Blaster worm and SoBig.F viruses of 2003 resulted in \$35 Billion in damages.¹ Since then, the magnitude of the security problem has increased significantly. In January 2007, Internet experts estimated that “botnet” programs – sophisticated programs that install themselves on unprotected personal computers – were present in more than 10 percent of the 650 million computers worldwide that are connected to the Internet. Botnet programs enable attackers to link infected computers into a powerful network that can be used to steal sensitive data, as well as money from online bank accounts and stock brokerages.²

While the software industry has made significant investments in writing more secure code, it is widely recognized that software vulnerability problems cannot be completely solved “ex-ante”; it is virtually impossible to design software that is free of vulnerabilities. Hence software firms continue to try to discover vulnerabilities after the software has been licensed.³ When vulnerabilities are identified “ex-post,” software firms typically issue updates (or patches) to eliminate the vulnerabilities. Those consumers who apply updates are protected in the event that attackers (or hackers) exploit the vulnerability.⁴ Applying updates is costly to consumers, however, and hence not all consumers necessarily apply them.⁵ For these consumers, the issuing of updates has a downside. The release of updates enables hackers to “reverse engineer” and find

¹ See “Internet security: Fighting the worms of mass destruction, *Economist*, Nov 27, 2003, available at http://www.economist.co.uk/science/displayStory.cfm?story_id=2246018.

² For example, one file created by a botnet program over a month contained about 55,000 login accounts (with passwords) and nearly 300 credit card numbers. Botnets also increase the damage caused by viruses because of their sophisticated, powerful communications network. See “Attack of the Zombie Computers is Growing Threat, John Markoff, *New York Times*, January 7, 2007, <http://www.nytimes.com/2007/01/07/technology/07net.html?em&ex=1168318800&en=79cc489d42f00bc8&ei=5087%0A>.

³ The intellectual property in software is typically “licensed” for use, not sold outright.

⁴ Granick (2005) remarks that “attacker” is the correct term, since hacker traditionally meant pioneer or explorer. However, the terms are now used interchangeably.

⁵ While some updates are automatically installed, other updates are not automatic and consumers must choose whether to install them. If it was indeed costless for consumers to install updates, there would not be any unprotected consumers. Meta Group Staff (2002) describes some of the costs consumers incur when they install updates.

10/24/2007

out how to exploit the vulnerabilities. The reverse engineering increases the probability of attack – and hence reduces the value of software to consumers who do not install updates.

The Slammer, Blaster, and Sobig.F viruses exploited vulnerabilities even though security updates had been released. That is, although the updates were widely available, relatively few users had applied them. Those consumers who did not install the updates suffered damages from these viruses. According to the *Economist*, the vulnerabilities exploited by these viruses were reverse engineered by hackers.⁶ Further, the time between the disclosure of a software vulnerability and the time in which an attack exploiting the vulnerability takes place has declined significantly. The Economist notes that the time from disclosure of the vulnerability to the time of attack was six months for the Slammer worm (January 2003), while the time from disclosure to attack for the Blaster worm (August 2003) was only three weeks.

There is a lively debate in the Law and Computer Science/Engineering literature about the pros and cons of disclosing vulnerabilities and the possibility of a regulatory regime requiring mandatory disclosure of vulnerabilities; see Swire (2004) and Granick (2005) for further discussion. Some security experts advocate full disclosure, in the belief that disclosure will provide incentives for software firms to make the software code more secure and to quickly fix vulnerabilities that are identified. Others advocate limited or no disclosure because they believe that disclosure significantly increases attacks by hackers. The debate is nicely summed up by Bruce Schneier, a well-known security expert: “If vulnerabilities are not published, then the vendors are slow (or don't bother) to fix them. But if the vulnerabilities are published, then hackers write exploits to take advantage of them.”⁷

It is not clear that it is possible to impose “mandatory disclosure” for vulnerabilities found by the firm who produces the software, since it can choose to keep the information to itself. But vulnerabilities are often discovered by third-parties and their policies can effectively impose mandatory disclosure. The Computer Emergency Response Team/Coordination Center (CERT/CC), for example, acts as an intermediary between those who report vulnerabilities and

⁶ See “Internet security: Fighting the worms of mass destruction, *Economist*, Nov 27, 2003, available at http://www.economist.co.uk/science/displayStory.cfm?story_id=2246018.

⁷ Schneier, B., “Crypto-Gram Newsletter,” February 15, 2000, available at <http://www.schneier.com/crypto-gram-0002.html>

10/24/2007

software vendors.⁸ When CERT/CC is notified about a potential vulnerability, it contacts the software vendor and gives it a 45 day period to develop a security update. It is CERT/CC's policy to then disclose the vulnerability even if a security update has not been made available by the firm. This policy essentially mandates disclosure of vulnerabilities that CERT/CC reports to the software vendors.⁹

When mandatory disclosure can be imposed, is it socially optimal to do so? Is CERT/CC policy welfare enhancing? What is the effect of disclosure policy on the price of the software, the market served, and firms' profits? How do reductions in the number of vulnerabilities and/or increases in the probability that the firm will find vulnerabilities before hackers affect disclosure policy, prices, profits, and welfare? In this paper, we develop a setting to examine the economic incentives facing software vendors and users when software is subject to vulnerabilities.¹⁰

We consider a firm that sells software which is subject to potential security breaches or vulnerabilities. The firm needs to set the price of the software and state whether it intends to disclose vulnerabilities and issue updates. Consumers differ in their value of the software and the potential damage that hackers may inflict on them. If the firm discloses vulnerabilities and provides updates, consumers who install updates are protected, even in the event that hackers exploit the vulnerability and attack, while consumers who do not install updates are worse off. Installing updates takes time and often requires re-booting systems. This is costly to consumers and they have to decide whether to install them.

The dilemma for the firm (regarding its disclosure policy) comes from the fact that the release of an update makes reverse engineering feasible for the hacker and increases the likelihood of attack. Such attacks cause damage to consumers who have not installed the updates. Thus, the

⁸ CERT/CC is a center for Internet security in the Software Engineering Institute at Carnegie Mellon University. Although CERT/CC is not formally a public agency, it acts as an intermediary between users and vendors.

⁹ CERT/CC is not the only source of vulnerabilities reported to software firms. Private security companies and benevolent users also identify software vulnerabilities and report them directly to software firms.

¹⁰ A recent paper by Polinsky and Shavell (2006) asks a similar question concerning product risks. In their model, the disclosure of product risk information is always beneficial to consumers and the benefit of voluntary disclosure arises from the firm's incentive to acquire more information about product risks because it can keep silent if the information is unfavorable. In our model, however, there is a third party (i.e., hackers) that can utilize the disclosed information to harm consumers. As a result, information disclosure can be harmful to consumers who do not update.

10/24/2007

decision of the firm to disclose and issue updates changes the value of software, increasing it for high-value users (who will employ updates when available) and decreasing it for low-value users (who will not employ updates when available). A third group of moderate-value users will install updates when available but indeed prefer a non-disclosure policy.

Since the availability of updates changes the value of the software, increasing it for some consumers and reducing it for others, the issuance of updates affects the firm's optimal price. Consequently, the firm's disclosure policy and its profit-maximizing behavior are interdependent. Our model derives the conditions under which a firm would disclose vulnerabilities. The firm's disclosure policy is not always socially optimal; hence we examine a regulatory policy that mandates disclosure of vulnerabilities. While a 'Mandatory Disclosure' regulatory policy is welfare improving in some cases, it is welfare reducing in other cases. This result sheds light on the source of the debate regarding a mandatory disclosure regulatory policy.

The firm can invest (ex-ante) to reduce the number of software vulnerabilities and/or invest ex-post to increase the probability that it will find problems before hackers. Reducing the number of potential vulnerabilities is equivalent to improving the quality of the software. Our model shows that ex-ante investment to reduce the number of vulnerabilities may lead to a "switch" from disclosure to a non-disclosure policy. Interestingly, such a regime switch can lead to a lower equilibrium price, despite the improvement in the quality of the software.

Ex-post investment increases the probability that the firm will find problems before hackers. When the firm optimally discloses vulnerabilities, such an increase raises profits and welfare. On the other hand, when the firm optimally does not disclose vulnerabilities, an increase in the probability of identifying them before hackers may induce the firm to switch to a disclosure policy and issue updates. This result sheds light on the effect of so-called "Bug Bounty" programs, in which firms offer rewards to users who identify and report vulnerabilities.¹¹

¹¹ Bug bounty programs have become quite popular and have attracted a lot of attention. In 2004 the Mozilla Foundation announced the Mozilla Security Bug Bounty program that rewards users who identify and report security vulnerabilities in the open source project's software. Under the program, users who report security bugs that are judged as critical by the Mozilla Foundation staff can collect a \$500 cash prize. See <http://www.mozilla.org/security/bug-bounty.html>. Independent security intelligence companies also offer a bounty for security bugs. TippingPoint, for instance, solicits hackers to report vulnerabilities in exchange for money under

Our paper builds on the nascent literature at the “intersection” of computer science/engineering and economics on cyber security. Much of the work in the field has been undertaken by computer scientists/engineers and legal scholars.¹² There is also a literature in management science that focuses on the tradeoff facing a software firm between an early release of a product with more security vulnerabilities and a later release with a more secure product.¹³ The few contributions by economists have focused on the lack of incentives for individuals or network operators to take adequate security precautions.¹⁴ Although the information security disclosure “dilemma” we examine in this paper is quite different, the economics literature has addressed the tradeoff between disclosure and non-disclosure in the context of intellectual property. In Anton and Yao (2004), for example, disclosure of intellectual property is beneficial because it enables a firm to receive a patent or to facilitate complementary innovation. But, disclosure is also costly since it enables imitation. In their setting, adopting a non-disclosure policy means the firm keeps a “trade-secret.”

The remainder of the paper is organized in the following way. Section 2 sets up the basic model of software market that is subject to potential security breaches. As a benchmark, we analyze the case in which the firm does not disclose vulnerabilities and there is no regulation requiring disclosure. Section 3 considers the case of mandatory disclosure regulation. In section 4, we analyze the firm’s voluntary incentives to disclose vulnerabilities. Section 5 investigates the effects of mandatory disclosure regulation on social welfare by comparing the market outcomes under voluntary and mandatory disclosure regimes. We consider the possibility of *ex ante* and *ex post* investments in reducing and identifying vulnerabilities in section 6, and analyze their effects

its “Zero Day Initiative” program. If a vulnerability is found, TippingPoint notifies the maker of the flawed product and updates its security products to protect users against exploitation of the flaw until an official update is released. IDefense, another security firm, recently offered \$10,000 to anyone who discovers a Windows flaw that leads to a critical fix under its “Vulnerability Contributor Program.”

¹² See Anderson (2006) for discussion.

¹³ See Arora, Caulkins, and Telang (forthcoming, 2007). August and Tunca (2006) consider a related setting. The key difference is that we allow for the possibility that a hacker will discover vulnerabilities before the firm. This is essential for the key tradeoff we examine, namely that disclosure endogenously increases the value of software to high value users and reduces it for low value users. August and Tunca (2006) also address different policy issues – they ask whether consumers should be required to install updates or whether they should be offered rebates for installing them.

¹⁴ This is because there is a “security” externality; individuals (or network operators) will not adequately protect against viruses on their computer (networks), since a large portion of the cost of the spread of the virus is incurred by others. See Varian (2004) and Camp and Wolfram (2004).

on the incentives to disclosure vulnerabilities and social welfare. Section 7 provides brief concluding remarks.

2. The Model

Consider a firm that produces a software product which is subject to potential security breaches or vulnerabilities. The number of expected security breaches is exogenously given and denoted by n .¹⁵ We assume that the firm is a sole producer of the software, we normalize production cost to zero, and we denote the price by p .

There is a continuum of consumers whose number is normalized to 1. Consumers are heterogeneous in terms of their valuation of the software and the damage incurred from an attack in the case of a security breach. We represent consumer heterogeneity by θ , assuming for convenience that θ is uniformly distributed on $[0,1]$.¹⁶ We assume that the value of software to consumer type θ is given by θv , where $v > 0$. Damage from each security breach exploited by hackers is assumed to be θD , where $D < v$. Hence, both the gross consumer valuation and the damage are increasing functions of consumer type. This assumption reflects the fact that while high valuation consumers benefit more from the software, they suffer more damage from an attack.

Consumers can either license (purchase) one unit of the software at the price p , or not purchase at all. Downloading and installing an update is costly to consumers; the cost is given by c , $c < D$. The cost (to consumers) of installing updates typically involves shutting the system down and restarting it, as well as possibly conducting some tests before installing the updates.¹⁷ As noted above, these actions take time and monetary resources.

After the product is sold, the firm continues to try to identify vulnerabilities. We assume that with probability α either the firm identifies the vulnerabilities itself before hackers, or

¹⁵ In section 6, we examine the effect of a reduction in the number of vulnerabilities on disclosure policy.

¹⁶ We assume a uniform distribution in order to derive closed-form solutions to our model. However, all the main qualitative results can be derived by assuming more general distributions with the monotone hazard rate property.

¹⁷ Firms typically do not charge consumers for updates.

institutions like CERT/CC, private security firms, or benevolent users find the vulnerabilities before hackers and report them to the firm. Thus, α is the percentage of problems that the firm finds or are reported to the firm by third-parties before they are discovered by hackers.¹⁸ When the firm discovers the security vulnerability before the hackers, it has an option to release an update, which protects those consumers who employ the update.

When hackers identify the security breach before the firm, all consumers who purchased the software are subject to potential damages. We do not explicitly model hacker preferences nor their decision making process. We simply assume that hackers attack with a fixed probability.¹⁹ We let $\gamma (< 1)$ be the probability that hackers will discover a vulnerability on their own (i.e., without disclosure) and attack. If the firm discloses the vulnerability and releases an update, we assume that the probability of attack is one. This assumption captures the fact that the release of an update makes reverse engineering feasible for the hacker and increases the likelihood of attack.

We consider three possible disclosure regimes:

- (i) The firm does not disclose any security vulnerability nor does it issue updates.
- (ii) The firm must disclose all security vulnerabilities and is obliged to release an update whenever it discovers a security vulnerability, or is informed about a vulnerability by a third party.
- (iii) The firm has the option of either adopting a policy to disclose vulnerabilities (and issue updates) or adopting a non-disclosure policy. The firm's disclosure policy is known to consumers at the time they purchase the software.

¹⁸ In the main part of the paper, α is given. In section 6 we examine the effect of an increase in the probability that the firm finds the security vulnerabilities before hackers on disclosure policy.

¹⁹ See Png, Tang, and Wang (2006) for an analysis that explicitly models hackers as a strategic player. They assume that hackers derive enjoyment from an attack on a user provided that they are not discovered by an enforcement agency. The focus of their paper is mainly on comparative statics results that analyze the direct and indirect effects of changes in the user cost of precaution and the rate of enforcement against hackers. Our focus, in contrast, is on software vendors' optimal decisions concerning voluntary disclosure and the effects of investment in security.

When the firm discloses vulnerabilities and issues updates, damage for a consumer who installs updates occurs only when hackers find the vulnerabilities before the firm finds them. Hence the net value to a consumer of type θ from purchasing the software and installing updates, denoted $W_u(\theta)$, is

$$(1) \quad W_u(\theta) = \theta v - \gamma(1-\alpha)\theta Dn - \alpha c n \equiv Z\theta - \alpha c n,$$

where $Z \equiv v - \gamma(1-\alpha)Dn$. $W_u(\theta)$ consists of the consumption value, the expected damage in the case where the hackers find the vulnerabilities before the firm, and the expected cost of installing updates. Similarly, $W_{nu}(\theta)$ is the net consumer value from buying the software, without installing updates.

$$(2) \quad W_{nu}(\theta) = \theta v - \gamma(1-\alpha)\theta Dn - \alpha\theta Dn \equiv S\theta,$$

where $S \equiv v - \gamma(1-\alpha)Dn - \alpha Dn$. The third term in $W_{nu}(\theta)$ is the expected damage to a consumer of type θ when the firm finds the security breach, discloses vulnerabilities, and issues an update which the consumer does not employ.

Finally, the value to a consumer of type θ from purchasing software when the firm does not disclose vulnerabilities, denoted $W_{nd}(\theta)$, is given by

$$(3) \quad W_{nd}(\theta) = \theta v - \gamma\theta Dn \equiv T\theta,$$

where $T \equiv v - \gamma Dn$. Comparing equations (1) - (3), yields $S < T < Z$. The differences among S , T , and Z are due to the differences in expected damage to consumers from an attack in these three cases.²⁰ We have $Z > T$, since a consumer of type θ who installs updates when the firm discloses vulnerabilities incurs less expected damage than in the case in which the firm does not disclose vulnerabilities; $T > S$, since the expected damage to a consumer of type θ who does not install updates is higher under a disclosure policy than under a non-disclosure policy because announcing vulnerabilities increases the probability of attack.

²⁰ The “damages” do not include the cost of installing updates.

We make the following two assumptions that hold throughout the paper:

- **A1:** We assume that $S > 0$, which guarantees that $W_{nu}(\theta) > 0$ for all θ .²¹ This assumption also implies that $W_u(\theta)$, $W_{nu}(\theta)$, and $W_{nd}(\theta)$ increase in consumer type θ .
- **A2:** We assume that $\gamma > c/D$. This assumption insures that $W_u(\theta) > W_{nd}(\theta)$ for at least some consumer types.

When A2 does not hold, i.e., when $\gamma < c/D$, the probability of a hacker attack is sufficiently small that software vulnerabilities are not a big concern. In such a case, the firm would never disclose vulnerabilities because $W_{nd}(\theta) > W_u(\theta)$ for every θ .

As a benchmark, we first consider case (i) in which the firm does not disclose vulnerabilities and there is no regulation that requires it to do so. Assumption (A1) guarantees that $W_{nd}(\theta)$ is increasing in θ . Hence given the firm's price, p , the consumers' purchase decision can be characterized by a threshold type $\theta_{nd}^*(p)$ such that only consumers of type $\theta \geq \theta_{nd}^*(p)$ will purchase the software. With the assumption of uniform distribution of θ , the number of buyers is given by $1 - \theta_{nd}^*(p)$.

Lemma 1: When the firm does not disclose vulnerabilities, the optimal price, market share, and profits are respectively given by $p_{nd}^* = T/2$, $(1 - \theta_{nd}^*) = 1/2$, and $\pi_{nd}^* = T/4$, where $T \equiv v - \gamma Dn$.²² ■

As intuition suggests, the profit-maximizing price and the firm's profits decrease in the probability of attack (γ), the number of vulnerabilities (n), and the damage (D) caused. When the firm does not disclose vulnerabilities, changes in α or c have no effect on the equilibrium price or profits.

3. The firm must disclose vulnerabilities and issue updates

²¹ This assumption is equivalent to assuming that all consumers will purchase the software at a zero price, regardless of whether they update or not.

²² For ease of presentation, all proofs are in the appendix.

Now consider a firm that is required to disclose identified vulnerabilities and issue an update that protects the software from these vulnerabilities. The firm cannot, however, require consumers to install updates.

In this setting equilibrium is defined as:

- A pricing strategy for the firm (p),
- A purchasing decision of a consumer type θ depending on the price and the number of software vulnerabilities.
- An updating decision of a consumer type θ , that specifies which consumers install updates.

Such that:

- (i) The price p is optimal given the consumers' purchasing and “update” behavior.
- (ii) The purchasing and “update” behavior are value maximizing for consumers.

We start by examining consumers' behavior. We first note that both $W_u(\theta)$ and $W_{nu}(\theta)$ are strictly increasing in θ by A1. In addition, $\frac{dW_u(\theta)}{d\theta} = Z > S = \frac{dW_{nu}(\theta)}{d\theta}$, which implies that $W_u(\theta)$ and $W_{nu}(\theta)$ cross at most once. Thus, we can characterize consumers' purchase and “update” behavior by two threshold types, $\theta^*(p)$ and $\hat{\theta}$, as stated in the following Lemma.

Lemma 2: There are two threshold levels: (i) $\hat{\theta} = c/D$ such that consumers of type $\theta \geq \hat{\theta}$, who purchase the software will install updates when they are available, while consumers with $\theta < \hat{\theta}$ will not install updates; (ii) Given a software price, p , there is a $\theta^*(p)$, such that only consumers of type $\theta \geq \theta^*(p)$ will purchase the software. ■

We can distinguish between two cases that are determined by the price that the firm charges. There is a critical price \hat{p} such that whenever $p < \hat{p}$, the resulting purchasing decision is such that $\theta^*(p) < \hat{\theta}$, while $p \geq \hat{p}$ results in purchasing decisions such that $\theta^*(p) \geq \hat{\theta}$, where $\hat{p} = \frac{cv}{D} - \gamma(1-\alpha)cn - \alpha cn$. Thus, when $p < \hat{p}$, there are three sets of consumers: 1- $\hat{\theta}$ consumers

purchase the software and apply updates, $\hat{\theta} - \theta^*(p)$ consumers purchase the software but do not apply updates, and $\theta^*(p)$ consumers do not purchase the software at all. (See Figure 1)

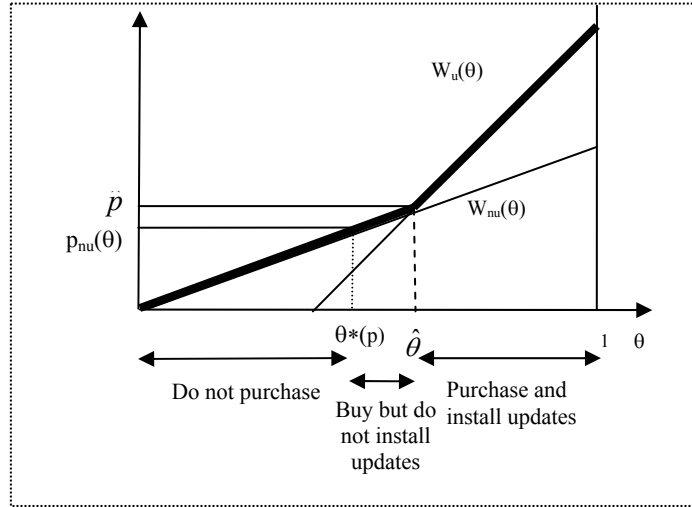


Figure 1: Purchase/Update Decision when Marginal Consumer Type $\theta^*(p) < \hat{\theta}$

It is more convenient to use θ as the firm's decision variable. For any θ , the price that the firm charges is defined by $p(\theta)$ which solves $\theta^*(p) = \theta$.²³ Whenever $\theta < \hat{\theta}$, the firm extracts the entire surplus from the marginal consumer θ who does not update. The software price in this case, denoted by $p_{nu}(\theta)$, satisfies the condition $p_{nu}(\theta) = \theta v - \gamma(1-\alpha)\theta Dn - \alpha\theta Dn$, and the firm's profit function is given by

$$(4) \quad \pi_{nu}(\theta) = p_{nu}(\theta) (1-\theta) = [\theta v - \gamma(1-\alpha)\theta Dn - \alpha\theta Dn](1-\theta) = S\theta(1-\theta).$$

The second case occurs whenever $p > \hat{p}$, which implies $\theta^*(p) > \hat{\theta}$; thus all the consumers who purchase the software will also install updates (see Figure 2). Since the marginal consumer installs updates, the software price satisfies the condition $p_u(\theta) = \theta v - \gamma(1-\alpha)\theta Dn - \alpha c n$, and the profits of the firm can be written:

$$(5) \quad \pi_u(\theta) = p_u(\theta) (1-\theta) = [\theta v - \gamma(1-\alpha)\theta Dn - \alpha c n](1-\theta) = (Z\theta - \alpha n c)(1-\theta).$$

²³ Since $\theta^*(p)$ is a strictly increasing function, the transformation is well defined.




Figure 2: Purchase/Updating Decision when Marginal Consumer Type $\theta^*(p) > \hat{\theta}$

The optimal pricing decision can be derived by solving $\text{Max} \{ \text{Max}_{\theta} \pi_{nu}(\theta); \text{Max}_{\theta} \pi_u(\theta) \}$.

Lemma 3: When the firm must disclose vulnerabilities and issues updates, the firm’s optimal price and profits are as follows:

(i) When $D/c < 2 - \alpha cn/Z$, the firm prefers to charge a low price and serve a larger market, including some consumers who do not install updates. The optimal price is $p_{nu}^* = S/2$, such that $p_{nu}^* < \hat{p}$; the number of consumers who purchase the software are $1 - \theta_{nu}^* = 1/2$, and the firm’s profits are $\pi_{nu}^* = S/4$.²⁴

(ii) When $D/c \geq 2 - \alpha cn/Z$, the firm will serve a smaller market of users, all of whom employ updates. The optimal price is $p_u^* = (Z - \alpha cn)/2$, such that $p_u^* > \hat{p}$; the number of consumers who purchase the software are $1 - \theta_u^* = (Z - \alpha cn)/2Z$, and the firm’s profits are $\pi_u^* = (Z - \alpha cn)^2 / (4Z)$.²⁵ ■

Intuitively, Lemma 3 shows that the firm’s optimal price and profits decrease with the number of vulnerabilities (n), the expected damage (D), and the probability of hacker attacks (γ) regardless of whether it sells only to consumers that update or to some consumers who do not update.

²⁴ Recall that $S \equiv v - \gamma(1 - \alpha)Dn - \alpha Dn$.

²⁵ Recall that $Z \equiv v - \gamma(1 - \alpha)Dn$. Since $D > c$, $Z > \alpha cn$ by assumption A1.

Increases in n , D , and γ make it more likely that the firm will serve a smaller market of high value consumers, all of whom install updates.²⁶

The effects of changes in α , the probability that the firm identifies the vulnerabilities before the hackers, on the firm's optimal prices and profits are more interesting. A higher α does not lead to an increase in software quality for all types of consumers. For consumers that do not install updates, a higher α implies a higher probability of hacker attack, and hence a lower software "quality." For consumers that install updates, a higher α means a more secure software program with lower expected damages, but consumers incur the cost of installing the updates.

Figure 3 illustrates the effect of increases in α on consumers' valuations $W_u(\theta)$ and $W_{nu}(\theta)$. Consumers that do not install updates are worse off and therefore $W_{nu}(\theta)$ goes down. For consumers who install updates, those with $\theta > c/D\gamma$ are better off and those with $\theta < c/D\gamma$ are worse off.²⁷ Consequently, the $W_u(\theta)$ curve rotates around the $\theta = c/D\gamma$ value.




Figure 3: Effects of an increase in α on $W_u(\theta)$, $W_{nu}(\theta)$

²⁶ Since π_u^* decreases in c and π_{nu}^* is independent of c , a decrease in c makes it more likely that the firm will serve a larger market, including some consumers who do not update.

²⁷ Assumption A2 insures that there are such types.

Proposition 1 (Effect of α , probability the firm identifies vulnerabilities before hackers):

(a) Suppose $D < (2-nc/v)c$. The marginal consumer does not install updates and the profit maximizing price and equilibrium profits decrease in α .

(b) Suppose $(2-nc/v)c \leq D < 2c$. There is a critical α , denoted $\hat{\alpha}(n, c, \gamma, D, v)$, such that when $\alpha \geq \hat{\alpha}$, the firm serves only consumers who install updates and when $\alpha < \hat{\alpha}$, the firm serves also some non-updaters.

(i) When α increases, but is still below $\hat{\alpha}$, the profit maximizing price and the equilibrium profits decrease in α .

(ii) The profit maximizing price increases discontinuously and the equilibrium market share falls discontinuously at $\alpha = \hat{\alpha}$.

(iii) When $\alpha \geq \hat{\alpha}$, an increase in α results in a higher price and a lower market share. Profits increase in α if and only if the probability of hacker attack is sufficiently large, i.e., if and only if $\gamma > \hat{\gamma}$, where $\hat{\gamma}$ is implicitly (and uniquely) defined by

$$\hat{\gamma} \equiv 2c / \left\{ \left[1 + \frac{\alpha n c}{Z(\hat{\gamma})} \right] D \right\}.^{28}$$

(c) Suppose $D \geq 2c$. The firm chooses to serve only consumers that install updates. Higher α results in a higher price and lower market share. Profits increase in α if and only if $\gamma > \hat{\gamma}$. ■

When D is relatively small or $\alpha < \hat{\alpha}$ (i.e., part (a) and part b(i) of Proposition 1), an increase in α decreases price and profits. This is because when D is relatively small or $\alpha < \hat{\alpha}$, the marginal consumer is a non-updater and the software becomes less valuable for the marginal user when α increases.

When D is relatively large or $\alpha > \hat{\alpha}$ (i.e., part b(iii) and part (c) of Proposition 1), the marginal consumer employs updates. In this case, a higher value of α increases the expected cost of installing updates, but also reduces the expected damages. The expected benefit exceeds the expected cost for consumer of types $\theta > c/D\gamma$, while the expected costs exceed the expected benefits for consumer of type $\theta < c/D\gamma$. An increase in α implies that the equilibrium price

²⁸ Hence $c/D < \hat{\gamma} < 2c/D$. It can be shown that $\hat{\gamma}$ decreases in α .

increases by $n(\gamma D - c)/2$, but the equilibrium market share falls by $Tnc/2Z^2$.²⁹ Thus, the affect of α on profits is not monotonic. Profits increase in α if and only if $\gamma > \hat{\gamma}$; when γ is large, the “higher price” effect dominates. When $\gamma < \hat{\gamma}$, the “lower market share” effect dominates and profits fall in α . We can conclude the following:

Corollary 1: When $D < (2-nc/v)c$, or $\alpha < \hat{\alpha}$, or $\gamma < \hat{\gamma}$, the firm’s optimal policy is to refrain from increasing α even when it is costless for the firm to do so and when it is costless to issue updates.

4. The Firm's Incentives to Disclose Vulnerabilities

Assume now that the firm has the option of choosing its disclosure policies. When the firm sells the software it can commit to disclosing vulnerabilities and issuing updates, or it can choose not to disclose vulnerabilities. The decision to disclose and issue updates affects the value of software. Figure 4 depicts the value of the software for consumers who do not install updates when available ($W_{nu}(\theta)$), for those who install updates when available ($W_u(\theta)$), as well as the value of software for the case in which the firm does not disclose vulnerabilities ($W_{nd}(\theta)$). A consumer that does not plan to install updates is always better off when the firm does not disclose vulnerabilities. In other words, the $W_{nu}(\theta)$ curve lies below the $W_{nd}(\theta)$ curve. Comparing $W_u(\theta)$ and $W_{nd}(\theta)$, there is a critical type, $\tilde{\theta} = c/D\gamma$, such that consumers of type $\theta > \tilde{\theta}$ are better off when the firm discloses vulnerabilities and consumers of type $\theta < \tilde{\theta}$ are better off when the firm does not disclose vulnerabilities. Note that $\tilde{\theta} > \hat{\theta}$ and “moderate-value” consumers of type $\theta \in [\hat{\theta}, \tilde{\theta}]$ will install updates when available, but prefer a non-disclosure policy.

As Figure 4 suggests, there are two possible outcomes when firms can set their disclosure policy: (i) the firm discloses vulnerabilities and sets a price such that $\theta^*(p) \geq \tilde{\theta}$ and all consumers install updates. (ii) the firm sets a price such that $\theta^*(p) < \tilde{\theta}$ and does not disclose vulnerabilities.

²⁹ $n(\gamma D - c)/2$ is greater than zero, since $\gamma > c/D$ by Assumption A2.




Figure 4: Willingness to pay under disclosure and non-disclosure

Proposition 2 (Disclosure Choice): The firm's optimal disclosure policy is to disclose vulnerabilities when $D\gamma/c \geq 2-\alpha cn/Z$ and not disclosure vulnerabilities when $D\gamma/c < 2-\alpha cn/Z$. ■

The condition in Proposition 2, $D\gamma/c \geq 2-\alpha cn/Z$, says that the firm will disclose vulnerabilities when the percentage gain in sales from doing so exceeds the percentage loss from lower prices. Since $\gamma < 1$, the condition, $D/c > 2-\alpha cn/Z$ from Lemma 3, holds whenever the condition from Proposition 2 holds. This means that when the firm discloses vulnerabilities it sells only to consumers that install updates.

Proposition 3 (Effect of the probability of hacker attack, γ , on Firm's Disclosure Policy):

There is a critical value of damage, $\tilde{D}(n, \alpha, c, v)$, such that

- (i) Whenever $D \leq \tilde{D}$, the firm will not disclose vulnerabilities.
- (ii) Whenever $D > \tilde{D}$, there is a critical probability of hacker attack, $\tilde{\gamma}(n, \alpha, D, c, v)$, such that whenever $\gamma \leq \tilde{\gamma}$, the firm will not disclose vulnerabilities, and whenever $\gamma > \tilde{\gamma}$, the firm discloses vulnerabilities.

\tilde{D} is defined implicitly (and uniquely) by $\tilde{D} = \left\{ 2 - \frac{\alpha cn}{[v - (1-\alpha)n\tilde{D}]} \right\} c$. ■

Proposition 3 shows that when the damage is relatively small, the firm will not disclose vulnerabilities, regardless of the value of γ . Whenever D is large, there is a critical probability of hacker attack, $\tilde{\gamma}(n, \alpha, D, c, v)$, such that when $\gamma > \tilde{\gamma}$, the firm discloses vulnerabilities.

Lemma 4:

- (i) $\tilde{\gamma}(n, \alpha, D, c, v)$ decreases in n and α .
- (ii) $\tilde{D}(n, \alpha, c, v)$ decreases in n and α . ■

5. Disclosure Policy, Regulation and Social Welfare

There is a debate among security experts regarding whether disclosure of software vulnerabilities should be mandatory. Some security experts recommend mandatory public disclosure of discoveries of potential security vulnerabilities, both to warn system administrators and users and to spur the vendor involved to develop an update as quickly as possible. Other experts are concerned that mandatory disclosure will lead to the reverse engineering (and exploitation) of vulnerabilities. As we discussed in the introduction, CERT/CC policy effectively mandates disclosure of vulnerabilities it reports to firms, while other regulations like the Digital Millennium Copyright Act can limit the publication of vulnerability information.³⁰ In this section, we examine the effect of a regulatory policy requiring disclosure on social welfare, i.e., we consider a regulator that can mandate the disclosure of vulnerabilities. Setting disclosure policy, however, does affect the market price as well as the number of consumers who purchase the software.

Since we assume no production costs, and since the price is a transfer from consumers to firms, social welfare is simply the integral of consumers' willingness to pay for software over the set of consumers who actually make the purchase. When the firm discloses vulnerabilities and $(D/c) < 2 - \alpha nc/Z$, the equilibrium is such that consumers of type $\theta \in [1/2, c/D]$ buy the software, but do not install updates, while consumers of type $\theta \in [c/D, 1]$, buy the software and install

³⁰ The Digital Millennium Copyright Act, which was primarily designed to protect intellectual property rights, has been used by the U.S. government and some firms to limit the publication of information about security vulnerabilities. See Granick (2005) for an expanded discussion.

updates. Summing up the surplus of these two groups of consumers gives us the total social surplus, denoted SW_{nu} ,³¹ in this case:

$$\begin{aligned}
 (6) \quad SW_{nu} &= \int_{1/2}^{c/D} W_{nu}(\theta) d\theta + \int_{c/D}^1 W_u(\theta) d\theta \\
 &= \int_{c/D}^1 \{[v - (1-\alpha)\gamma Dn]\theta - \alpha nc\} d\theta + \int_{1/2}^{c/D} [v - (1-\alpha)\gamma Dn - \alpha Dn]\theta d\theta \\
 &= \frac{3Z}{8} - \alpha Dn \frac{(4c^2 - D^2)}{8D^2} - \alpha nc \frac{(D-c)}{D}.
 \end{aligned}$$

When the firm discloses vulnerabilities and $(D/c) > 2-\alpha nc/Z$, the equilibrium is such that the firm sells only to consumers of type $\theta \in [\frac{(Z + \alpha nc)}{2Z}, 1]$ (See Lemma 3). Since these consumers also install updates, the total social surplus, denoted SW_u , is:

$$(7) \quad SW_u = \int_{1/2 + \alpha nc / 2Z}^1 W_u(\theta) d\theta = \int_{1/2 + \alpha nc / 2Z}^1 \{[v - (1-\alpha)\gamma Dn]\theta - \alpha nc\} d\theta = \frac{3Z}{8} + 3 \frac{(\alpha nc)^2}{8Z} - 3\alpha nc / 4.$$

Finally, when the firm adopts a non-disclosure policy, the equilibrium is such that it sells to consumers of type $\theta \in [1/2, 1]$. Total social surplus in this case, denoted SW_{nd} , is

$$(8) \quad SW_{nd} = \int_{1/2}^1 W_{nd}(\theta) d\theta = \int_{1/2}^1 (v - \gamma Dn)\theta d\theta = \frac{3T}{8}.$$

The regulator adopts the disclosure policy that maximizes social welfare as defined by (6)-(8). Figure 5 shows the firm's optimal disclosure policy and the regulator's disclosure policy as a function of the parameters of the model.

³¹ The subscript "nu" signifies the fact that the marginal consumer who is indifferent between purchase and no purchase "does not update."




Figure 5: Regulator vs. Market Outcome

Figure 5 shows that, depending on the parameters, there are five possible regions:

- Region I: Suboptimal Disclosure (Firm does not Disclose; Regulator would Disclose.)
- Region II: Efficient (Firm does not Disclose; Regulator would not Disclose.)
- Region III: Efficient (Firm does not Disclose; Regulator would not Disclose.)
- Region IV: Efficient (Firm does Disclose; Regulator would Disclose.)
- Region V: Efficient (Assumption A2 does not hold; hence neither the firm nor the regulator would Disclose.)

Proposition 4 (Regulator vs. Market Outcome): The equilibrium disclosure policy of the firm is socially optimal unless the parameters are in Region I (Figure 5), in which case mandatory disclosure is optimal whereas the firm prefers not to disclose. Region I is bounded by two conditions, which are $\beta < 2 - \alpha nc/Z$ and $\gamma > (8\beta - 4 - \beta^2)/3\beta^2 \equiv \bar{\gamma}$, where $1 < \beta \equiv D/c < 2$. ■

In Region I, the firm will choose not to disclose vulnerabilities, while welfare maximization requires such a disclosure. The divergence between the firm and the regulator is because the regulator's disclosure policy depends on the effect of disclosure on the *average* consumer, whereas the vendor's profit-maximizing disclosure policy depends on the impact on the *marginal* consumer. Since there are heterogeneous consumers, the average consumer type cares more about security than the marginal type. This effect leads to suboptimal disclosure in the market in Region I. Although the "average/marginal" effect exists in Region II as well, the probability of hacker attack is sufficiently low in this region so that neither the firm nor the regulator would disclose vulnerabilities.

In Regions III and IV, there is a second effect that offsets the “average/marginal consumer” effect. The opposing effect is that market share is higher under a non-disclosure regime. A regulator values an increased market share more than the firm does, because the firm obtains the full surplus only from the marginal consumer. In our setting, these opposing effects exactly cancel out. Thus in Regions III and IV, the market outcome is efficient: A regulator would mandate disclosure whenever the firm would disclose vulnerabilities.³²

Corollary 2: Mandatory disclosure increases social welfare in Region I, but reduces welfare in Regions II and III. In Region IV, Mandatory Disclosure has no effect, since the firm discloses vulnerabilities. ■

Corollary 2 illustrates the source of the debate regarding a mandatory disclosure regulatory policy. Mandatory disclosure is welfare improving in one region, but welfare reducing in other regions. Mandatory disclosure also affects equilibrium prices, as well as the number of consumers that purchase the software.

Corollary 3 (The Effect of Mandatory Disclosure on Equilibrium Prices):

- (i) In Regions I and II in Figure 5, mandatory disclosure decreases the equilibrium price.
- (ii) In Region III, mandatory disclosure increases the equilibrium price and reduces equilibrium number of consumers.
- (iii) In Region IV, mandatory disclosure has no effect on either the price or the number of consumers who purchase software. ■

In Regions I and II, the firm would not disclose vulnerabilities in the absence of regulation. Since the marginal user is a non-updater under disclosure, mandatory disclosure lowers the willingness to pay for the marginal consumer; hence it will lead to a lower equilibrium price. In Region III, the firm would not disclose vulnerabilities in the absence of regulation. Since all

³² If, for example, θ was not uniformly distributed, the two effects present in Regions III and IV would not cancel out and the inefficiency (suboptimal or excess disclosure) would depend on the distribution of consumer types. But this would not change the main result of Corollary 2 (below) that mandatory disclosure can be welfare reducing as well as welfare improving.

consumers install updates under mandatory disclosure in this case, the firm serves a smaller market of higher quality-sensitive consumers. Hence, in this case, mandatory disclosure leads to a higher equilibrium price but reduces the firm's market share. In Region IV, the firm indeed discloses vulnerabilities in the absence of regulation. Hence, mandatory disclosure has no effect in this case.

6. Ex-Ante and Ex-Post Investment in Reducing and Identifying Security Vulnerabilities

There are two types of investments the firm can undertake: (i) Investment that reduces the number of software vulnerabilities (i.e., reducing n) and (ii) Investment that increases the probability that the firm will find the software vulnerabilities before hackers (i.e., increasing α). The first type of investment can be thought of as an ex-ante investment in quality, while the second type can be thought of as an ex-post investment.

6.1 Ex-Ante Investment to Reduce the Number of Software Vulnerabilities

Many software firms now provide formal training in order to teach their programmers how to write code that is less vulnerable to attacks.³³ This can be interpreted as an investment in reducing the number of software vulnerabilities before the software is sold. A reduction in n , hereafter denoted as Δn , can be viewed as an increase in the quality of the product for all consumer types; thus it raises consumer willingness to pay for the software (See Figure 6).

³³ “Several initiatives are underway to improve secure programming skills and knowledge. Symantec, Oracle, Microsoft, and a few other software companies are conducting short courses for their programmers; software firms like SPI Dynamics and Fortify Technology are working with universities to provide automated, real-time feedback to student programmers; and dozens of universities are creating elective courses on secure programming.” (quote taken from <http://www.sans-ssi.org/#pgoals>.) Additionally, the SysAdmin, Audit, Network, Security (SANS) Software Security Institute recently launched a new initiative involving more than 360 companies, government agencies and colleges to help software developers, programmers and students improve their knowledge of how to write secure software code. The press release of the initiative can be found at http://www.sans-ssi.org/ssi_press.pdf.




Figure 6: Effects of a decrease in n on $W_u(\theta)$, $W_{nu}(\theta)$, $W_{nd}(\theta)$

Proposition 5 (Ex-Ante investment):³⁴ Suppose $\tilde{\gamma}(n, \alpha) < \gamma < 2c/D$ and $D > \tilde{D}(n, \alpha)$. If Δn is sufficiently large so that $\gamma < \tilde{\gamma}(n - \Delta n, \alpha)$ or $D < \tilde{D}(n - \Delta n, \alpha)$, the reduction Δn will induce a switch from a disclosure policy to a non-disclosure policy. Such a switch may be accompanied by a lower equilibrium price. Otherwise, a reduction in n leads to an increase in the equilibrium price, profits and consumers welfare, but has no effect on the disclosure policy of the firm. ■

Although a reduction in n is an improvement in software quality, the higher quality does not necessarily imply a higher equilibrium price. The intuition is that when $\tilde{\gamma}(n, \alpha) < \gamma < 2c/D$ and $D > \tilde{D}(n, \alpha)$ the firm's optimal policy is to disclose vulnerabilities. Since both $\tilde{\gamma}(n, \alpha)$ and $\tilde{D}(n, \alpha)$ are decreasing functions of n , a reduction in n results in a higher $\tilde{\gamma}(n, \alpha)$ and a higher $\tilde{D}(n, \alpha)$. If the reduction is sufficiently large such that $\gamma < \tilde{\gamma}(n - \Delta n, \alpha)$ or $D < \tilde{D}(n - \Delta n, \alpha)$, it induces the firm to switch from a disclosure policy to a non-disclosure policy. Such a regime change is welfare improving, despite the fact that it induces the firm not to disclose

³⁴ The parameters of interest here are α and n . Hence, we write $\tilde{\gamma}(n, \alpha)$ and $\tilde{D}(n, \alpha)$ rather than $\tilde{\gamma}(n, \alpha, D, c, v)$ and $\tilde{D}(n, \alpha, c, v)$.

vulnerabilities. Moreover, it establishes a non-monotonicity (and possible discontinuity) in the software price as a function of software quality (n), which is caused by a switch in the disclosure policy of the firm.

Corollary 4: There is a critical level, denoted \tilde{n} ; $\tilde{n} = \frac{v(2c - \gamma D)}{(1 - \alpha)\gamma D(2c - \gamma D) + \alpha c^2}$, such that a regime change (from disclosure to non-disclosure) occurs when $n > \tilde{n}$ and $n - \Delta n < \tilde{n}$. The price of software will fall under the regime change if and only if $\Delta n/n < \alpha(\gamma D - c)/\gamma D$. ■

6.2 Ex-Post Investment: Increasing α

Assume that the firm can increase the probability that it finds vulnerabilities before the hackers find them or that third-party policies increase α . In Proposition 1 and Figure 3, we considered the effect of higher α on prices and profits in the case in which the firm was required to disclose vulnerabilities. In such a case, a higher α may reduce prices and profits. We now extend the analysis and consider the effect of a higher α on the firm's disclosure policy, and well as on prices, profits, and welfare.

Proposition 6 (Ex-Post investment):

- (i) When $\gamma > \tilde{\gamma}(n, \alpha)$ and $D > \tilde{D}(n, \alpha)$, the firm would disclose vulnerabilities and an increase in α implies a higher price, greater profits, and higher welfare without any change in the firm's disclosure policy.
- (ii) When $\gamma < \tilde{\gamma}(n, \alpha)$ or $D < \tilde{D}(n, \alpha)$, the firm does not disclose vulnerabilities regardless of the value of D . A relatively small increase in α does not change disclosure policy and does not affect the price or firm profits. A relatively large increase in α may induce the firm to adopt a policy of disclosure; a change in disclosure policy results in a higher price, greater profits, and higher welfare. ■

In case (i), the firm discloses vulnerabilities since $\gamma > \tilde{\gamma}(n, \alpha)$ (Proposition 3). Furthermore, $\partial \tilde{\gamma}(n, \alpha)/\partial \alpha < 0$ and $\partial \tilde{D}(n, \alpha)/\partial \alpha < 0$ and thus $\gamma > \tilde{\gamma}(n, \alpha)$ implies $\gamma > \tilde{\gamma}(n, \alpha + \Delta \alpha)$ and $D > \tilde{D}(n, \alpha + \Delta \alpha)$. Consequently, disclosure is optimal regardless of the magnitude of the increase

in α . Profits increase in α in this case because $\tilde{\gamma}(n, \alpha) > \hat{\gamma}(n, \alpha)$. Since $SW_u = 3\pi_u^*/2$, an increase in profits increases Social Welfare as well.

In case (ii), the optimal policy is not to disclose vulnerabilities. But since $\tilde{\gamma}(n, \alpha)$ and $\tilde{D}(n, \alpha)$ are decreasing functions of α , an increase in α results in a lower $\tilde{\gamma}(n, \alpha)$ and a lower $\tilde{D}(n, \alpha)$. If the increase in α is relatively small, the firm continues not to disclose vulnerabilities. Since π_{nd}^* is independent of α , the equilibrium price and profits are unchanged. On the other hand, a large increase in α may induce a switch from case (ii) to case (i). A switch from a non-disclosure policy to a disclosure policy takes place if $\gamma > \tilde{\gamma}(n, \alpha + \Delta\alpha)$ and $D > \tilde{D}(n, \alpha + \Delta\alpha)$.

Proposition 6 shows that when the firm can choose its disclosure policy, ex-post investment either leads to higher prices, greater profits, and higher welfare or does not affect prices, profits or welfare. Thus, unlike the case when the firm is required to disclose vulnerabilities, when the firm can choose its disclosure policy, a higher α never reduces prices and profits. Proposition 6(ii) shows that a higher α may also induce the firm to make a (welfare-improving) shift from non-disclosure to disclosure.

Proposition 6 has interesting implications for the effects of “Bug Bounty” programs, in which firms (or third parties) offer rewards to users who identify and report vulnerabilities. The introduction of a bounty program, in which vulnerabilities “bought” through the program by third parties are provided to firms, can be interpreted in our setting as an increase in α .³⁵ Proposition 6(i) implies that the use of a bounty program has a positive effect on both profitability and welfare. This is because in such a case (Region IV in Figure 5), the firm discloses vulnerabilities, the marginal consumer applies updates, and profits and welfare are increasing in α . In case (ii), the introduction of a bounty program has no effect if, despite the increase in α , the firm continues to employ a non-disclosure policy (Region III in Figure 5).³⁶ If the increase in α is large enough, however, the introduction of a bounty program will induce the

³⁵ We assume the bounty program, if offered by independent security companies, is such that the vulnerability will be disclosed only when an update is available from software vendors.

³⁶ Although the firm who supplies the software would not introduce a ‘Bounty Program’ here, a third party might do so.

firm to switch to case (i), i.e., from non-disclosure to a disclosure policy (or from Region III to Region IV in Figure 5). This is because the boundary between regions III and IV in Figure 5 shifts down and to the left when α increases.

7. Concluding Remarks and Further Discussion

In this paper, we examined the incentives for a software firm to adopt a disclosure or non-disclosure policy and the interdependence between the pricing of software and the disclosure policy. We used our framework to examine public policies suggested by security experts: Mandatory Disclosure of vulnerabilities and Bug Bounty programs. We find that Mandatory Disclosure is not necessarily welfare improving. Mandatory disclosure improves welfare only when the probability of attack is very high and the expected damage is relatively small. When both the probability of attack and the expected damage are moderate, Mandatory Disclosure is welfare reducing, since a non-disclosure policy maximizes welfare. When both the probability of attack and damage are large, Mandatory Disclosure has no effect since the firm would disclose vulnerabilities even without regulatory intervention. We find that a Bug Bounty program is a welfare improving policy instrument since it either has no effect on the firm's disclosure policy or it induces a welfare-improving change in disclosure policy (from non-disclosure to disclosure).

Finally, we considered the possibility that the firm could invest in identifying software vulnerabilities. The investment can be either ex-ante investment in which the number of vulnerabilities is reduced prior to the release (and sale) of the software, or ex-post investment in which the firm increases the probability that it will identify vulnerabilities ex-post before hackers do so and issues the appropriate updates. An ex-ante reduction in the number of vulnerabilities typically leads to higher prices, greater profits, and higher welfare, but it may also induce a (welfare-improving) regime shift from a disclosure to non-disclosure policy. Such a regime shift may be accompanied by lower prices, despite the increase in software quality. Ex-post investment also typically leads to higher prices, greater profits, and higher welfare, but it also may induce a (welfare-improving) regime shift in the opposite direction: from non-disclosure to disclosure.

10/24/2007

References

- American Online and the National Cyber Security Alliance, *AOL/NCSA Online Safety Study*, October 2004.
- Anderson, R., and T. Moore, 2006, "The Economics of Information Security," *Science*, 314:610-613
- Anton, J., and D. Yao, 2004, "Little Patents and Big Secrets: Managing Intellectual Property," *Rand Journal of Economics*, 35:1-22.
- Arora, A., Caulkins, J.P., and R. Telang, "Sell First, Fix Later: Impact of Patching on Software Quality," *Management Science*, forthcoming, 2007.
- August, T., and T. Tunca, 2006, "Network Software Security and User Incentives," *Management Science*, 52:1703-1720.
- Camp, L.J., and C. Wolfram, "Pricing Security," in L.J. Camp and S. Lewis, eds., *Economics of Information Security*, vol. 12, *Advances in Information Security*. Springer-Kluwer, 2004.
- Granick, J., 2005, "The Price of Restricting Vulnerability Publications," *International Journal of Communications Law and Privacy*, 9: 1-35.
- Meta Group Staff, "META Report: Security Vulnerability Disclosures," January 2002, available at http://itmanagement.earthweb.com/it_res/article.php/947271
- Png, Ivan, Tang, Qian, and Wang, Qihong, "Information Security: Use Precautions and Hacker Targeting," 2006, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=912161
- Polinsky, A. M., and S. Shavell, "Mandatory versus Voluntary Disclosure of Product Risks," Harvard John M. Olin Discussion Paper No. 564, October 2006.
- Schneier, B., 2000, "Crypto-Gram Newsletter," available at <http://www.schneier.com/crypto-gram-0002.html>
- Swire, P., 2004, "A Model for When Disclosure Helps Security: What is Different about Computer and Network Security, Theory of Software for Security and Competitive Reasons: Open Source, Proprietary Software," *Journal on Telecommunications and High Technology Law*, 163.
- Varian, H., 2004, "System Reliability and Free Riding," available at <http://www.ischool.berkeley.edu/~hal/Papers/2004/reliability>.

Appendix: Proof of Lemmas, Propositions, and Corollaries

Proof of Lemma 1:

There is one-to-one correspondence between the price and the marginal consumer type who is indifferent between purchasing and no purchasing. It is more convenient to use the marginal type (θ_{nd}) as the firm's choice variable. Since the firm captures the surplus of the marginal consumer, the price and profits are as follows:

$$p_{nd}(\theta_{nd}) = [\theta_{nd}v] - \gamma\theta_{nd}Dn \equiv T\theta_{nd}$$

$$\pi_{nd}(\theta_{nd}) = p_{nd}(\theta_{nd})(1-\theta_{nd}) = [\theta_{nd}v - \gamma\theta_{nd}Dn](1-\theta_{nd}) \equiv T\theta_{nd}(1-\theta_{nd})$$

Maximizing these profits yields $p_{nd}^* = T/2$, $1-\theta_{nd}^* = 1/2$, and $\pi_{nd}^* = T/4$. ■

Proof of Lemma 2:

(i) Comparing $W_u(\theta)$ and $W_{nu}(\theta)$ yields a threshold consumer, $\hat{\theta}$, where $\hat{\theta} = c/D$. In addition, $W_u(\theta)$ is steeper than $W_{nu}(\theta)$ since we have $\frac{dW_u(\theta)}{d\theta} = Z > S = \frac{dW_{nu}(\theta)}{d\theta}$. Therefore, $W_u(\theta) \leq$

$W_{nu}(\theta)$ for all $\theta \leq \hat{\theta}$ and $W_u(\theta) \geq W_{nu}(\theta)$ for all $\theta \geq \hat{\theta}$.

(ii) Since both $W_u(\theta)$ and $W_{nu}(\theta)$ are increasing in θ , the function $\text{Max}\{W_u(\theta), W_{nu}(\theta)\}$ is also increasing in θ and therefore, given a price p , there is a marginal consumer type, denoted $\theta^*(p)$, such that only consumers of type $\theta \geq \theta^*(p)$ will purchase the software. Given our assumption of a uniform distribution of types, $1-\theta^*(p)$ is the number of consumers who purchase the software and $\theta^{*\prime}(p) \geq 0$. ■

Proof of Lemma 3:

Note that $\pi_{nu}(\theta) = p_{nu}(\theta)(1-\theta) = [\theta v - \gamma(1-\alpha)\theta Dn - \alpha\theta Dn](1-\theta) = S\theta(1-\theta)$. Thus, $\pi_{nu}(\theta)$ is maximized when $\theta_{nu}^* = 1/2$ with the optimal price of $p_{nu}^* = S/2$, which yields the profit of $\pi_{nu}^* = S/4$. In contrast, $\pi_u(\theta) = p_u(\theta)(1-\theta) = [\theta v - \gamma(1-\alpha)\theta Dn - \alpha nc](1-\theta) = (Z\theta - \alpha nc)(1-\theta)$. It can be easily verified that $\pi_u(\theta)$ is maximized when $\theta_u^* = (Z + \alpha nc)/2Z$ with the optimal price of $p_u^* = (Z - \alpha nc)/2$. The number of consumers who purchase the software is $1-\theta_u^* = (Z - \alpha nc)/2Z$. The maximum profit is given by $\pi_u^* = (Z - \alpha nc)^2/(4Z)$. By noticing that $S = Z - \alpha Dn$, we can easily verify that $\pi_{nu}^* > \pi_u^*$ if and only if $D/c < 2-\alpha nc/Z$. In addition, we can verify that $\theta_{nu}^* = 1/2 < \hat{\theta}$ when $D/c < 2-\alpha nc/Z$ with $\pi_{nu}^* > \pi_u^*$, which proves claim (i). Similarly, $\theta_u^* = (Z + \alpha nc)/2Z > \hat{\theta}$ when $D/c > 2-\alpha nc/Z$ with $\pi_{nu}^* < \pi_u^*$, which proves claim (ii). ■

Proof of Proposition 1:

(a) For ease of presentation in the proofs, we define Condition (C1) as $D/c < 2-\alpha nc/Z$, which is the condition for $\pi_{nu}^* > \pi_u^*$ in Lemma 3; we see that the RHS of Condition (C1) decreases in α while the LHS of Condition (C1) does not depend on α . Hence, the RHS is minimized with the value of $2-nc/v$, since $Z=v$, when $\alpha=1$. Thus, Condition (C1) always holds if $D/c < 2-nc/v$. Rewriting, Condition (C1) always holds when $D < (2-nc/v)c$. When $D < (2-nc/v)c$ the marginal consumer does not install updates (by Lemma 3) and $p_{nu}^* = S/2 = [v - \gamma(1-\alpha)Dn - \alpha Dn]/2$. Hence

$\partial p^{nu} / \partial \alpha = - (1-\gamma)Dn/2 < 0$, since $\gamma < 1$. $\pi_{nu}^* = [v - \gamma(1-\alpha)Dn - \alpha Dn]/4 = [v - \gamma Dn - \alpha(1-\gamma)Dn]/4$.
Hence $\partial \pi^{nu} / \partial \alpha = - (1-\gamma)Dn/4 < 0$.

(b) Condition (C1) can be rewritten $\alpha n(D-2c) + \frac{(\alpha cn)^2}{Z} < 0$. Since the second term is greater than zero, Condition (C1) does not hold when $D \geq 2c$. By continuity, when $(2-nc/v)c \leq D < 2c$ there exists a unique $\hat{\alpha}$ such that condition (C1) holds if and only if $\alpha < \hat{\alpha}$, where $\hat{\alpha}$ is implicitly defined by $D/c = 2 - \hat{\alpha} nc/Z(\hat{\alpha})$.

(i) $\alpha < \hat{\alpha}$ (and $\alpha + \Delta\alpha < \hat{\alpha}$): the marginal consumer does not update, the result follows from (a).

(ii) $\alpha < \hat{\alpha}$ and $\alpha + \Delta\alpha > \hat{\alpha}$: there is a regime change and this causes the discontinuity.

(iii) $\alpha > \hat{\alpha}$: the marginal consumer installs updates.

$\partial p_u^* / \partial \alpha = \{\gamma Dn - nc\}/2 > 0$, since $\gamma > c/D$ by assumption. $\partial(1-\theta^u) / \partial \alpha = -Tnc/2Z^2 < 0$.

$$\pi_u^* = \frac{(Z - \alpha nc)^2}{4Z} = \frac{1}{4} \left\{ Z - 2\alpha nc + \frac{(\alpha nc)^2}{Z} \right\}. \quad \partial \pi_u^* / \partial \alpha = \frac{1}{4} \left\{ \gamma Dn - 2nc + \alpha n^2 c^2 \frac{(2Z - \gamma \alpha Dn)}{Z^2} \right\}.$$

Since the third term is greater than zero, $\gamma > 2c/D$ is a sufficient condition for profits to increase in α . We now find a sufficient and necessary condition: Let $\gamma = \delta c/D$.

$$\partial \pi_u^* / \partial \alpha = \frac{1}{4} \left\{ \delta cn - 2nc + \alpha n^2 c^2 \frac{(2Z - \delta \alpha cn)}{Z^2} \right\} = \frac{cn}{4} \left\{ \delta - 2 + \frac{\alpha nc}{Z} \left[2 - \delta \frac{\alpha cn}{Z} \right] \right\} =$$

$$\frac{cn}{4} \left\{ \delta - 2 + x(2 - \delta x) \right\}, \text{ where } x = \alpha nc/Z. \text{ Since } Z > \alpha nc, x < 1. \partial \pi_u^* / \partial \alpha > 0 \Leftrightarrow$$

$$\left\{ \delta - 2 + x[2 - \delta x] \right\} > 0 \Leftrightarrow \delta(1-x^2) > 2(1-x) \Leftrightarrow \delta > 2/(1+x).$$

Thus, the sufficient and necessary condition for $\partial \pi_u^* / \partial \alpha > 0$ can be written as $\gamma > \Psi(\gamma)$, where

$$\Psi(\gamma) = \frac{2c}{\left(1 + \frac{\alpha nc}{Z}\right)D} \text{ (note that } Z \text{ is a function of } \gamma). \text{ It can be easily verified that } \Psi(\gamma) \text{ is a}$$

strictly decreasing function of γ . In addition, when $\alpha > \hat{\alpha}$ and thus $D/c > 2 - \alpha nc/Z$ holds, we can show that $\Psi(\gamma) < 1$ since $\alpha nc/Z < 1$. We also know that $\Psi(\gamma=0) > 0$. Taken together, this implies that there is a unique $\hat{\gamma}$ that is implicitly defined by $\hat{\gamma} \equiv \frac{2c}{\left(1 + \frac{\alpha nc}{Z}\right)D}$ ($\hat{\gamma}$ is implicitly defined since

Z is a function of γ), such that:

$$\gamma > \hat{\gamma} \Rightarrow \partial \pi_u^* / \partial \alpha > 0$$

$$\gamma = \hat{\gamma} \Rightarrow \partial \pi_u^* / \partial \alpha = 0$$

$$\gamma < \hat{\gamma} \Rightarrow \partial \pi_u^* / \partial \alpha < 0$$

(c) $D \geq 2c$: the marginal consumer installs updates. The results follow from (b). ■

Proof of Proposition 2:

The proof parallels that of Lemma 3. More specifically, algebraic manipulation shows that $\pi^{nd} > \pi^l$ if and only if $\gamma D/c < 2 - \alpha nc/Z$ (which we refer to as condition (C2) for ease of presentation in the proofs of other propositions that follow). ■

Proof of Proposition 3:

The LHS of Condition (C2) increases in γ , while the RHS decreases in γ . When $\gamma=1$, the LHS is still smaller than the RHS whenever $D < [2 - \alpha nc / (v - (1 - \alpha)Dn)]c \equiv \Omega(D)$. Then, we can find a unique \tilde{D} , which is implicitly defined by $\tilde{D} = \Omega(\tilde{D})$, such that $D < [2 - \alpha nc / (v - (1 - \alpha)Dn)]c$ if and only if $D < \tilde{D}$. Thus, when $D < \tilde{D}$, the firm will not disclose vulnerabilities regardless of the value of γ . By continuity, whenever $D > \tilde{D}$, there exists a $\tilde{\gamma}$ such that the firm is indifferent between disclosing and not disclosing vulnerabilities. When $\gamma < \tilde{\gamma}$, Condition (C2) holds and the firm will not disclose vulnerabilities. When $\gamma > \tilde{\gamma}$, Condition (C2) does not hold and the firm will disclose vulnerabilities. ■

Proof of Lemma 4:

(i) We first show that $\tilde{\gamma}(n, \alpha)$ is a decreasing function of n . From Condition (C2), $\tilde{\gamma}$ is implicitly

defined by: $\tilde{\gamma} = \frac{2c}{D} - \frac{\alpha nc^2}{Z(\tilde{\gamma})D}$. We let $k \equiv \frac{\alpha c^2}{D}$, and thus

$$\frac{\partial \tilde{\gamma}}{\partial n} = -k \frac{\partial}{\partial n} \left[\frac{n}{Z(\tilde{\gamma}, n)} \right] = -k \frac{Z - n \left[\frac{\partial Z}{\partial \tilde{\gamma}} \cdot \frac{\partial \tilde{\gamma}}{\partial n} + \frac{\partial Z}{\partial n} \right]}{Z^2}.$$

By rearranging terms, we can rewrite the equation above as

$$\frac{\partial \tilde{\gamma}}{\partial n} \left[Z^2 - kn \cdot \frac{\partial Z}{\partial \tilde{\gamma}} \right] = -kZ + kn \cdot \frac{\partial Z}{\partial n}$$

Note that $\frac{\partial Z}{\partial \tilde{\gamma}} = -(1 - \alpha)Dn < 0$ and $\frac{\partial Z}{\partial n} = -\tilde{\gamma}(1 - \alpha)Dn < 0$. Therefore,

$$\frac{\partial \tilde{\gamma}}{\partial n} = \frac{-kZ + kn \cdot \frac{\partial Z}{\partial n}}{Z^2 - kn \cdot \frac{\partial Z}{\partial \tilde{\gamma}}} < 0$$

Now, we show that $\frac{\partial \tilde{\gamma}}{\partial \alpha} < 0$:

$\tilde{\gamma}$ is implicitly defined by: $\tilde{\gamma} = \frac{2c}{D} - \frac{\alpha nc^2}{Z(\tilde{\gamma})D}$, we denote $\omega \equiv \frac{nc^2}{D}$, and thus

$$\frac{\partial \tilde{\gamma}}{\partial \alpha} = -k \cdot \frac{\partial}{\partial \alpha} \left[\frac{\alpha}{Z(\tilde{\gamma}, \alpha)} \right] = -k \frac{Z - \alpha \left[\frac{\partial Z}{\partial \tilde{\gamma}} \cdot \frac{\partial \tilde{\gamma}}{\partial \alpha} + \frac{\partial Z}{\partial \alpha} \right]}{Z^2}$$

$$\frac{\partial \tilde{\gamma}}{\partial \alpha} \left[Z^2 - k\alpha \cdot \frac{\partial Z}{\partial \tilde{\gamma}} \right] = -kZ + k\alpha \frac{\partial Z}{\partial \alpha}$$

$$\frac{\partial \tilde{\gamma}}{\partial \alpha} \left[Z^2 + k\alpha(1 - \alpha)Dn \right] = -kZ + k\alpha \tilde{\gamma} Dn$$

$$\frac{\partial \tilde{\gamma}}{\partial \alpha} \underbrace{[Z^2 + n^2 c^2 (1 - \alpha)]}_{>0} = -kT < 0$$

$$\Rightarrow \frac{\partial \tilde{\gamma}}{\partial \alpha} < 0.$$

We show that $\tilde{D}(n, \alpha)$ is a decreasing function of n . \tilde{D} is uniquely defined by $\tilde{D} = 2c - \frac{\alpha n c^2}{Z(\tilde{D}, n)}$.

Therefore, we have

$$\frac{\partial \tilde{D}}{\partial n} = -\alpha c^2 \cdot \frac{\partial}{\partial n} \left[\frac{n}{Z(\tilde{D}, n)} \right] = -\alpha c^2 \cdot \frac{Z - n \left[\frac{\partial Z}{\partial \tilde{D}} \cdot \frac{\partial \tilde{D}}{\partial n} + \frac{\partial Z}{\partial n} \right]}{Z^2}, \text{ which can be rewritten as}$$

$$\frac{\partial \tilde{D}}{\partial n} [Z^2 - \alpha n c^2 \cdot \frac{\partial Z}{\partial \tilde{D}}] = -\alpha c^2 [Z - n \frac{\partial Z}{\partial n}]. \text{ Since } \frac{\partial Z}{\partial \tilde{D}} = -\gamma(1 - \alpha)n < 0, \text{ we have } \frac{\partial \tilde{D}}{\partial n} < 0.$$

We now show that $\tilde{D}(n, \alpha)$ is a decreasing function of α .

$$\text{Again, } \tilde{D} = 2c - \frac{\alpha n c^2}{Z(\tilde{D}, \alpha)}$$

$$\frac{\partial \tilde{D}}{\partial \alpha} = -n c^2 \cdot \frac{\partial}{\partial \alpha} \left[\frac{\alpha}{Z(\tilde{D}, \alpha)} \right] = -n c^2 \cdot \frac{Z - \alpha \left[\frac{\partial Z}{\partial \tilde{D}} \cdot \frac{\partial \tilde{D}}{\partial \alpha} + \frac{\partial Z}{\partial \alpha} \right]}{Z^2}$$

$$\frac{\partial \tilde{D}}{\partial \alpha} [Z^2 - \alpha n c^2 \cdot \frac{\partial Z}{\partial \tilde{D}}] = -n c^2 \underbrace{[Z - \alpha \cdot \gamma D n]}_{=T}$$

$$\frac{\partial \tilde{D}}{\partial \alpha} \underbrace{[Z^2 - \alpha n c^2 (-\gamma(1 - \alpha)n)]}_{>0} = -T n c^2 < 0$$

$$\Rightarrow \frac{\partial \tilde{D}}{\partial \alpha} < 0.$$

■

Proof of Proposition 4:

By using equations (6) and (8) in the main text, we know that $SW_{nu} > SW_{nd}$ if and only if

$$\frac{3(v - \gamma D n)}{8} + \frac{3\gamma \alpha D n}{8} - \alpha D n \left\{ \frac{4c^2 - D^2}{8D^2} - \alpha n c \frac{(D - c)}{D} \right\} > \frac{3(v - \gamma D n)}{8}. \text{ By denoting } \beta = D/c, \text{ this}$$

condition can be rewritten as $\gamma > \frac{8cD - 4c^2 - D^2}{3D^2} = \frac{8\beta - 4 - \beta^2}{3\beta^2} \equiv \bar{\gamma}$. We also know that if $D/c < 2 -$

$\alpha n c / Z$, we have $\pi_{nu}^* > \pi_u^*$ (see Lemma 3), which implies that the firm will choose a price that induces the marginal consumer not to update if the firm is mandated to disclose vulnerability. Proposition 2, however, tells us that the firm has no incentive to disclose if $D\gamma/c < 2 - \alpha n c / Z$, which always holds when $D/c < 2 - \alpha n c / Z$. Hence Region I in which the firm would not disclose, but a regulator would obtain if $\gamma > \bar{\gamma}$ and $D/c < 2 - \alpha n c / Z$.

To see the alignment between private and social incentives to disclose in other regions, note that

$$SW_{nd} = \frac{3(v - \gamma Dn)}{8} = 3 \frac{T}{8} = \frac{3\pi_{nd}^*}{2}$$

$$\pi_u^* = \frac{(Z - \alpha nc)^2}{4Z} = \frac{Z}{4} - \alpha nc/2 + \frac{(\alpha nc)^2}{4Z}$$

$$SW_u = \frac{3[v - \gamma(1 - \alpha)Dn]}{8} - 3\alpha nc/4 + \frac{3(\alpha nc)^2}{8Z} = 3 \frac{Z}{8} - 3\alpha nc/4 + \frac{3(\alpha nc)^2}{8Z} = \frac{3\pi_u^*}{2}$$

Hence $SW_u > SW_{nd}$ iff $\pi_u^* > \pi_{nd}^*$. ■

Proof of Corollary 2:

The Region I result follows immediately from the proof of proposition 4.

In the case of Region II, we have $\pi_{nd}^* > \pi_{nu}^* > \pi_u^*$ and the firm has no incentive to disclose. It also implies that if there is a mandatory disclosure regulation, the firm will choose a price that results in a marginal consumer who does not choose apply updates. In Region II, we also have $SW_{nu} < SW_{nd}$ since $\gamma < \bar{\gamma}$. Thus, mandatory disclosure regulation in this region would lower social welfare than the original outcome of no disclosure.

In the case of Region III, we have $\pi_{nd}^* > \pi_u^* > \pi_{nu}^*$, which implies that mandatory disclosure results in the marginal consumer applying updates. However, we also know that $SW_{nd} > SW_u$ in this region. Therefore, once again, mandatory mandatory disclosure regulation in this region would lower social welfare than the original outcome of no disclosure.

In the case of Region IV, mandatory disclosure does not make any difference since the market outcome is disclosure.

Proof of Corollary 3:

(i) In regions I and II, Condition (C1) holds and thus mandatory disclosure changes the regime from non-disclosure to disclosure where the marginal consumer does not update. This causes a change of price from p_{nd}^* to p_{nu}^* , where $p_{nu}^* < p_{nd}^*$ (from Lemmas 1 and 3, and since $T > S$). $1 - \theta_{nu}^* = 1 - \theta_{nd}^* = 1/2$ (from Lemmas 1 and 3), and thus the equilibrium market share is unaffected.

(ii) In region III Condition (C1) does not hold and Condition (C2) holds; thus mandatory disclosure changes the regime from non-disclosure to disclosure with the marginal consumer updating. This causes a change of price from p_{nd}^* to p_u^* , from Lemmas 1 and 3: $p_u^* = (Z - \alpha nc)/2$, $p_{nd}^* = T/2$, which implies that $p_u^* > p_{nd}^*$ since $Z - \alpha cn > T$ (from the definitions of Z , T and assumption A2). $1 - \theta_u^* = 1/2 - \alpha cn/2Z < 1/2 = 1 - \theta_{nd}^*$, so the equilibrium market share decreases.

(iii) In region IV conditions C1 and C2 do not hold, thus mandatory disclosure does not cause a regime change. The equilibrium remains “disclosure” and all consumers install updates. The price and market share remain p_u^* and $1 - \theta_u^*$. ■

Proof of Proposition 5:

By Proposition 3, we know that if $\tilde{\gamma}(n, \alpha) < \gamma < 2c/D$ and $D > \tilde{D}(n, \alpha)$, the firm does not disclose vulnerabilities. In addition, from Corollary 1, we know that $\tilde{\gamma}(n, \alpha)$ and $\tilde{D}(n, \alpha)$ decrease in n . Thus, if Δn is sufficiently large, we could have a situation of $\gamma < \tilde{\gamma}(n - \Delta n, \alpha)$ or $D < \tilde{D}(n - \Delta n, \alpha)$, triggering a switch to a non-disclosure regime.

For a small Δn there will be no regime change because of continuity. When we are in a disclosure regime, small reductions in n result in a higher price since $\frac{\partial p^u}{\partial n} = \frac{-\gamma(1-\alpha)D - \alpha c}{2} < 0$. The result would also be higher profits for the firm (and higher consumer welfare) since $\tilde{\gamma} > \hat{\gamma}$ (see proof of Proposition 6).

Note that $\gamma > 2c/D$ is a sufficient condition for the firm to disclose vulnerabilities, regardless of n . When the condition holds, a reduction in n increases prices and profits:

$$\frac{\partial \pi_u^*}{\partial n} = \frac{-(Z + \alpha nc)(Z - \alpha nc)\gamma(1 - \alpha)D - 2Z(Z - \alpha nc)\alpha c}{4Z^2} < 0$$

$$\frac{\partial p_u^*}{\partial n} = \frac{-\gamma(1 - \alpha)D - \alpha c}{2} < 0$$

When $\gamma < \tilde{\gamma}(n, \alpha)$, the firm finds it optimal not to disclose vulnerabilities. Since $\tilde{\gamma}(n, \alpha)$ is a decreasing function of n , a reductions in n do not affect the disclosure policy. In such a case, a reduction in n increases prices and profits:

$$\frac{\partial \pi_{nd}^*}{\partial n} = \frac{-\gamma D}{4}$$

$$\frac{\partial p_{nd}^*}{\partial n} = \frac{-\gamma D}{2}$$

Proof of Corollary 4:

When $\tilde{\gamma}(n, \alpha) < \gamma < 2c/D$, there exists a critical value of n , denoted \tilde{n} , for which $\pi_u^* = \pi_{nd}^*$. The condition $\pi_u^* = \pi_{nd}^*$ can be rewritten as $Z(\gamma D - 2c) + \alpha nc^2 = 0$. By solving the equation, we can derive the critical value of n as $\tilde{n} = \frac{v(2c - \gamma D)}{(1 - \alpha)\gamma D(2c - \gamma D) + \alpha c^2}$.

When $n > \tilde{n}$, we have $\pi_u^* > \pi_{nd}^*$ and thus the firm chooses disclosure.

When $n < \tilde{n}$, we have $\pi_u^* < \pi_{nd}^*$ and thus the firm chooses non-disclosure.

Let n be initial value of the number of vulnerabilities. Consider a situation in which $n > \tilde{n}$, but $n - \Delta n < \tilde{n}$, where Δn denotes the decline in the number of vulnerabilities. In such a case, the initial software price is given by $p_u^* = [Z(n) - \alpha nc]/2$ whereas the post-change price with $n - \Delta n$ ($< \tilde{n}$) vulnerabilities is given by $p_{nd}^* = [T(n - \Delta n)]/2$. The condition for $p_{nd}^* < p_u^*$ can be rewritten as $v - \gamma D(n - \Delta n) < v - \gamma(1 - \alpha)Dn - \alpha cn$, which yields the desired condition $\Delta n/n < \frac{\alpha(\gamma D - c)}{\gamma D}$.

Proof of Proposition 6:

The claims (i) and (ii) follow from Propositions 1 and 3 if $\tilde{\gamma} > \hat{\gamma}$ because $\frac{\partial \tilde{\gamma}}{\partial \alpha} < 0$ and $\frac{\partial \tilde{D}}{\partial \alpha} < 0$ (Corollary 1). Hence, we prove the Proposition by proving the condition ($\tilde{\gamma} > \hat{\gamma}$):

From the equation $\pi_u^* = \pi_{nd}^*$ we have $\tilde{\gamma}$ implicitly defined by $\tilde{\gamma} = \frac{2c}{D} - \frac{\alpha nc^2}{Z(\tilde{\gamma})D}$. Remember that $\hat{\gamma}$

is defined by $\hat{\gamma} = \frac{2Z(\hat{\gamma})c}{(Z(\hat{\gamma}) + \alpha nc)D}$ (see Proposition 1).

Let us denote the functions on the RHS of the equations that implicitly define $\tilde{\gamma}$ and $\hat{\gamma}$ as $\Omega(\gamma)$ and $\Psi(\gamma)$:

$$\Omega(\gamma) = \frac{2c}{D} - \frac{\alpha nc^2}{Z(\gamma)D}, \quad \Psi(\gamma) = \frac{2Z(\gamma)c}{(Z(\gamma) + \alpha nc)D}$$

Note that both $\Omega(\gamma)$ and $\Psi(\gamma)$ are strictly decreasing functions of γ . Therefore, a sufficient condition for $\tilde{\gamma} > \hat{\gamma}$ is $\Omega(\gamma) > \Psi(\gamma)$ for $\forall 0 < \gamma < 1$, that is,:

$$\frac{2c}{D} - \frac{\alpha nc^2}{Z(\gamma)D} > \frac{2Z(\gamma)c}{(Z(\gamma) + \alpha nc)D} \quad \forall 0 < \gamma < 1.$$

By cross-multiplying both sides of the inequality and collecting terms, we derive an equivalent condition as:

$$2Z(\gamma)Z(\gamma) + 2Z(\gamma)\alpha nc - \alpha ncZ(\gamma) - (\alpha nc)^2 > 2Z(\gamma)Z(\gamma)$$

Canceling out terms and dividing by αnc , we have $Z(\gamma) - \alpha nc > 0$, which is satisfied by our assumption A1. ■