Network Security: Vulnerabilities and Disclosure Policy[#]

by

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

January 29, 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 the consumers who install updates, while the disclosure itself facilitates reverse engineering of the vulnerability by hackers. The paper develops a setting that examines the economic incentives facing software vendors and users when software is subject to vulnerabilities. We consider a firm that sells software which is subject to potential security breaches. 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 and need to decide whether to purchase the software as well as whether to install updates. Prices, market shares, and profits depend on the disclosure policy of the firm. The paper analyzes the market outcome and derives the conditions under which a firm would disclose vulnerabilities. It then examines the effect of a regulatory policy that requires mandatory disclosure of vulnerabilities. The paper discusses the incentives to invest in product security by investigating how a decline in the number of vulnerabilities and an increase in the probability that the firm will identify vulnerabilities ex-post (before hackers) affect disclosure policy, price and profits.

JEL Classification: L100, L630.
Keywords: Internet security, software vulnerabilities, disclosure policy.

* Department of Economics, Michigan State University, 101 Marshall Hall, East Lansing, Michigan 48824-1038, Tel: 517-353-7281, E-mail: choijay@msu.edu

** The Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel, Tel: 972-3-640-7167, E-mail: fersht@post.tau.ac.il

*** Department of Public Policy, Tel Aviv University, Tel Aviv 69978, Israel, Tel: 972-3-640-6742, E-mail: gandal@post.tau.ac.il

[#] We are grateful to Sagit Bar-Gill for excellent research assistance and thank Jacques Lawarree, Shlomit Wagman, and participants from the WEIS 2005 and DIMACS 2007 conferences for their helpful comments. A research grant from Microsoft is gratefully acknowledged. Any opinions expressed are those of the authors.

1. Introduction

The Internet provides many benefits, but at the same time 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 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. 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.[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. For these consumers, the issuing of updates has a downside. The release of updates to eliminate vulnerabilities enables hackers to

[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] 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.

"reverse engineer" and find out how to exploit the vulnerabilities.[5] This 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.

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, market share, and profits. Consequently, the firm's disclosure policy and its profit-maximizing behavior are interdependent. In some cases it will be optimal for the firm to commit to supply updates, even though such updates are typically provided free of charge to consumers. In other cases it will be optimal for the firm to refrain from providing updates, even when the updates are without cost to the firm.

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 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

[5] In this context, reserve engineering is detrimental. For a detailed discussion of benefits from reverse engineering in the context of innovation, see Samuelson and Scotchmer (2005).

[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.

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.[8] 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 software vendors.[9] 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.[10] 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.[11]

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? In this paper, we develop a setting to examine the economic incentives facing software vendors and users when software is subject to vulnerabilities.

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. They need to decide whether to purchase the software as well as whether to install updates. If the firm discloses vulnerabilities and provides

[7] Schneier, B., "Crypto-Gram Newsletter," February 15, 2000, available at http://www.schneier.com/crypto-gram-0002.html

[8] A U.S. federal law enacted in 2005, however, requires U.S. financial institutions to notify their customers when a database storing their personal information has been breached.

[9] 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.

[10] CERT/CC first investigates to determine whether a security vulnerability indeed exists.

[11] 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.

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. Thus the firm's disclosure policy affects consumers' willingness to pay for the software.

Installing updates is costly to consumers and they themselves have to decide whether to install them. Not all consumers will necessarily choose to install updates. The dilemma for the firm comes from the fact that the release of an update makes reverse engineering feasible for the hacker and increases the likelihood of attack. Disclosure makes it easier for hackers to engage in a damaging activity and such attacks cause damage to consumers who have not installed the updates.

Our model derives the conditions under which a firm would disclose vulnerabilities. We show that prices are higher when the firm chooses to disclose vulnerabilities, while the firm serves a larger market when it does not disclose vulnerabilities. Disclosure of vulnerabilities is not always optimal for the firm. Even when it is costless for the firm to disclose vulnerabilities and issue updates, the firm will not necessarily choose to do so.

The firm's disclosure policy is not always socially optimal; hence we examine a regulatory policy that mandates disclosure of vulnerabilities. Such a policy is problematic, however, since in some circumstances non-disclosure is socially optimal. We identify two opposing effects that determine whether a firm has "suboptimal" or "excessive" incentives to disclose vulnerabilities.

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 in reducing 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. But when the firm optimally discloses vulnerabilities, such an increase raises prices and profits. 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.

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."

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

[12] See Anderson (2006) for discussion.

[13] See Arora, Caulkins, and Telang (forthcoming, 2007).

[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).

[15] In section 7, we examine the effect of a reduction in the number of vulnerabilities on disclosure policy.

in the case of a security breach. We represent consumer heterogeneity by a parameter θ, assuming for convenience that θ is uniformly distributed on [0,1]. 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)[16] 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, $c<D$.[17] The cost of installing updates typically involves shutting the system down and restarting it, as well as possibly conducting tests before installing the updates. These actions take time and monetary resources.[18]

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 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.[19]

When the firm discovers the security vulnerability before the hackers and releases an update, only those consumers who do not employ an update are unprotected. When hackers identify the security breach before the firm, there is no update and 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. We letγ, $\gamma<1$, be the probability that hackers will discover a vulnerability on their own (i.e., without disclosure) and attack. If the firm

[16] Although software is typically licensed, for ease of presentation, we will use the term "purchase."

[17] Firms typically do not charge consumers for updates.

[18] See Meta Group Staff (2002).

[19] In section 7 we examine the effect of an increase the probability that the firm finds the security vulnerabilities before hackers on disclosure policy.

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. This is equivalent to assuming that disclosure leads to an increase in expected damages for consumers who do not install updates.

We consider three possible disclosure regimes:

(i) 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.

(ii) The firm does not disclose any security vulnerability nor does it issue updates.

(iii) The firm can either adopt a policy to disclose vulnerabilities (and issue updates) or adopt a non-disclosure policy. The firm's disclosure policy is known to consumers at the time they purchase the software.

When the firm adopts a disclosure policy and issues updates, damage for a consumer who installs updates occurs only when the hacker finds the vulnerabilities before the firm finds them. Hence the net consumer's value from buying the software and installing an update, denoted $W_u(\theta)$, is

$$(1) \qquad W_u(\theta) = \theta v - \gamma(1-\alpha)\theta Dn - \alpha nc \equiv Z\theta - \alpha nc,$$

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

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

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

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

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] 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.

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

When A2 does not hold, the probability of a hacker attack is sufficiently small that software vulnerabilities are not a concern. When $\gamma<c/D$, the firm's optimal policy is non-disclosure of vulnerabilities, i.e., it would never disclose vulnerabilities and issue updates.

3. The firm must disclose vulnerabilities and issue updates

We now consider case (i) in which the firm is required to disclose identified vulnerabilities and it must issue an update that protects the software from these vulnerabilities. The firm cannot, however, require consumers to install updates.

[20] The "damages" do not include the cost of installing updates.

In this setting equilibrium is defined as:

- p - A pricing strategy for the firm,
- $B(\theta,p,n)$ - A purchasing decision of a consumer type θ depending on the price and the number of software vulnerabilities, where $B(\theta,p,n)=1$ if the consumer purchases the software and $B(\theta,p,n)=0$ if the consumer does not purchase it.
- $I(\theta) \in \{0,1\}$ - A decision of a consumer type θ, where $I(\theta)=1$ ($I(\theta)=0$) means that he/she does (does not) install the update.

Such that:

(i) The price p is optimal given the consumers' purchasing and "update" strategy.

(ii) $B(\theta,p,n)$ and $I(\theta)$ are value maximizing behavior of consumers.

Comparing $W_u(\theta)$ and $W_{nu}(\theta)$ yields a threshold consumer, $\hat{\theta}$, where $\hat{\theta}=c/D$, so that $W_u(\theta) \geq W_{nu}(\theta)$ for all $\theta \geq \hat{\theta}$. Thus, consumers of type θ, $\theta \geq \hat{\theta}$, who purchased the software will install updates when they are available, while consumers with $\theta < \hat{\theta}$ do not install updates.

Since both $W_u(\theta)$ and $W_{nu}(\theta)$ are increasing in θ, the function 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 $\theta'(p) \geq 0$.

We can distinguish between two cases that are determined endogenously by the price that the firm charges. There is a critical price p^* such that whenever $p<p^*$, the resulting purchasing decision is such that $\theta(p)<\hat{\theta}$, while $p \geq p^*$ results in purchasing decisions such that $\theta(p) \geq \hat{\theta}$. When $p<p^*$, 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. 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 $\theta(p)= \theta$. (See Figure 1)

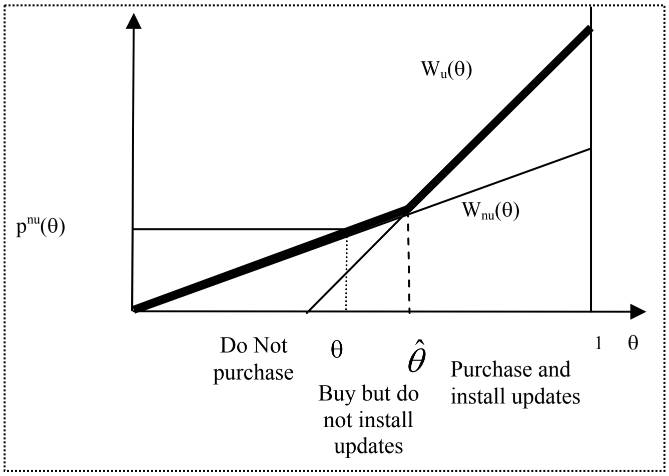

Figure 1: Purchase/Update Decision when Marginal Consumer less than $\hat{\theta}$

Whenever $\theta<\hat{\theta}$**,** the firm extracts the entire surplus from the marginal consumer θ who does not update. The software price, 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

$$(4) \qquad \pi^{nu}(\theta) = p^{nu}(\theta)\ (1-\theta)=\{\theta v - \gamma(1-\alpha)\theta Dn- \alpha\theta Dn\}(1-\theta)=S\theta(1-\theta)$$

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

$$(5) \qquad \pi^u(\theta) = p^u(\theta)\ (1-\theta) = \{\theta v - \gamma(1-\alpha)\theta Dn - \alpha nc\}(1-\theta)=(Z\theta- \alpha nc)(1-\theta)$$

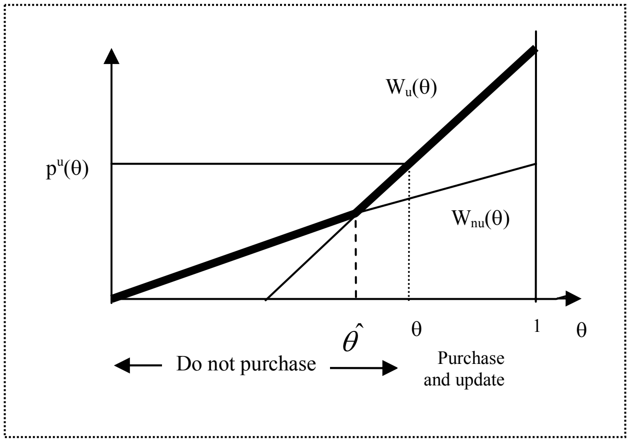

Figure 2: Purchase/Patching Decision when Marginal Consumer greater than $\hat{\theta}$

The pricing decision of the firm can be described as follows. Using (4) and (5), the firm needs to separately determine the optimal θ in the $[0,\hat{\theta}]$ and $[\hat{\theta},1]$ ranges and then needs to solve Max{Max$_\theta$ $\pi^{nu}(\theta)$, Max_θ $\pi^u(\theta)$}. The solution of this maximization problem yields the optimal price.

Proposition 1: When the firm must disclose vulnerabilities and issues updates, the optimal price and the firm's profits are as follows:[21]

(i) When Condition (C1) holds, the optimal price is $p^{nu} < p^*$ and is given by $p^{nu} = S/2$; the number of consumers who purchase the software are $1\text{-}\theta^{nu} = ½$, and the firm's profits are $\pi^{nu} = S/4$.[22]

(ii) When condition (C1) does not hold, the optimal price is $p^u > p^*$ and is given $p^u = (Z - \alpha nc)/2$; the number of consumers who purchase the software are $1\text{-}\theta^u = (Z - \alpha nc)/2Z$, and the firm's profits are $\pi^u = (Z - \alpha nc)^2/(4Z)$.[23]

Condition (C1) is given by $(1/2)(\alpha n(D-c)/2) < (\alpha nc/2Z)(Z-\alpha nc)/2$. ■

The intuition for Condition (C1) is as follows: The firm faces a tradeoff between serving a larger market with a lower price versus a smaller market with a higher price. From Proposition 1, the

[21] For ease of presentation, all proofs are in the appendix.

[22] Recall that $S \equiv v - \gamma(1\text{-}\alpha)Dn\text{-} \alpha Dn$.

[23] Recall that $Z \equiv v - \gamma(1\text{-}\alpha)Dn$. Since $D>c$, $Z>\alpha nc$ by assumption A1.

equilibrium price is lower by $\alpha n(D\text{-}c)/2 > 0$ when Condition (C1) holds. The equilibrium market share, on the other hand, is larger by $\alpha nc/2Z$ when Condition (C1) holds. Hence, the left hand side of condition (C1) is the equilibrium market share when the marginal consumer does not update ($1\text{-}\theta^{nu} = 1/2$) multiplied by the difference in equilibrium prices ($\alpha n(D\text{-}c)/2$), while the right hand side is the difference in market shares ($\alpha nc/2Z$) multiplied by the equilibrium price when the firm serves a smaller market ($[Z\text{-}\alpha nc]/2$). We can write Condition (C1) as $\Delta P/P^u < \Delta Q/Q^{nu}$. Hence Condition (C1) says that the firm will serve a larger market (i.e., sell to both those who update and those who do not update) when the percentage gain in market share from doing so exceeds the percentage loss in price.

Proposition 1 implies intuitively that the profit maximizing price and profits decrease with number of vulnerabilities (*n*), the expected damage (*D*), and the probability of hacker attacks (γ) regardless of whether condition (C1) holds or does not hold.[24]

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 is more interesting. Does the firm benefit when it identifies a larger percentage of the vulnerabilities before hackers? Do consumers necessarily benefit from it? That is not always the case. We first state the following Lemma.

Lemma 1:

(i) Suppose that $D<(2\text{-}nc/v)c$. Condition (C1) holds regardless of the value of α.

(ii) Suppose that $D\geq 2c$. Condition (C1) does not hold regardless of the value of α.

(iii) Suppose that $(2\text{-}nc/v)c\leq D<2c$. There is a critical α, denoted $\alpha^c(n,c,\gamma,D,v)$, such that when $\alpha<\alpha^c$, Condition (C1) holds and when $\alpha>\alpha^c$, Condition (C1) does not hold. ■

Lemma 1 (i) shows that when D is small relative to the cost of installing updates, Condition (C1) holds regardless of the value of α. Lemma 1 (ii) shows that when *D* is relatively large, Condition (C1) does not hold regardless of the value of α. Lemma 1(iii) shows that when *D* falls in

[24] When condition (C1) holds, prices and profits are unaffected by changes in the cost of the installing updates. When condition (C1) does not hold, equilibrium prices and profits fall when the cost of updates rises. Hence an increase in *c* makes it more likely that the Condition (C1) will hold.

intermediate range, condition (C1) holds for relatively small values of α. When α is increased above a certain threshold level, Condition (C1) ceases to hold and the firm will find it optimal to serve a smaller market and charge a higher price. In such a case a higher α may have a considerable effect on the market equilibrium. Thus in analyzing the effect of higher α on the equilibrium outcomes, we will distinguish between these cases.

Proposition 2 **(Effect of α):**

(a) Suppose $D<(2-nc/v)c$. The profit maximizing price and equilibrium profits decrease in α.

(b) Suppose $(2-nc/v)c\leq D<2c$.

(i) When α increases, but is still below α^c, the profit maximizing price and equilibrium profits decrease in α.

(ii) When the initial α was such that condition (C1) holds, i.e., $\alpha<\alpha^c$, but $\alpha>\alpha^c$ following the increase in α, the profit maximizing price increases discontinuously while the equilibrium market share falls discontinuously at $\alpha=\alpha^c$.

(iii) When $\alpha>\alpha^c$, 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 defined by $\hat{\gamma}\equiv 2c/\{[1+\frac{\alpha nc}{Z(\hat{\gamma})}]D\}$.[25]

(c) Suppose $D\geq 2c$.

An increase in α results in a higher price and a lower market share. A higher α implies greater profits if and only if $\gamma>\hat{\gamma}$. ■

Proposition 2 shows that an increase in α may decrease profits. That is, increasing the probability that the firm identifies the vulnerabilities ex-post and issues updates is not necessarily beneficial to the firm, even in situations when it is costless to the firm to increase α and costless to issue updates. This is because consumers do not necessarily benefit from an increase in α.

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

To understand the implications of changes in α we first depict the effect of a higher α on consumers valuations $W_u(\theta)$, $W_{nu}(\theta)$. Consumers that do not install updates are worse off and therefore $W_{nu}(\theta)$ goes down. For consumers who install updates, those with $\theta > c/D\gamma$ [26] are better off and those with $\theta < c/D\gamma$ are worse off. Consequently, the $W_u(\theta)$ curve rotates around the $\theta = c/D\gamma$ value. (See Figure 3)

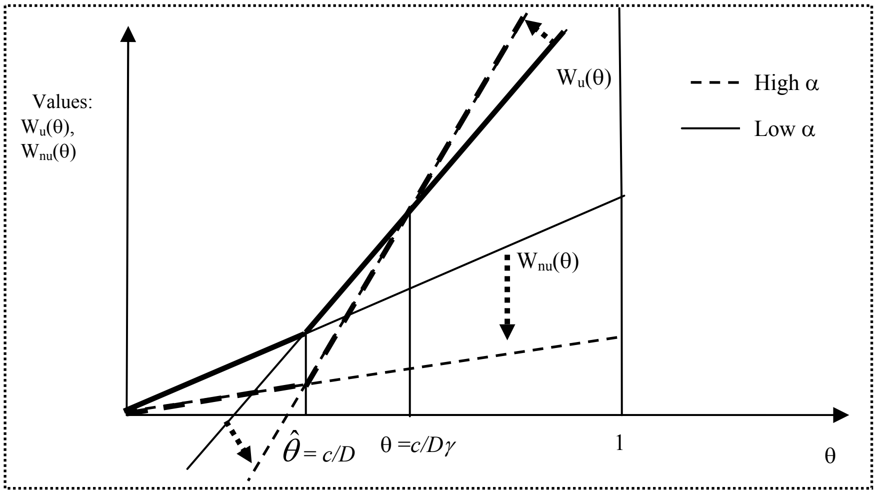

Figure 3: Effects of an increase in α on $W_u(\theta)$, $W_{nu}(\theta)$

The intuition for part a and part b(i) of Proposition 2 is that when Condition (C1) holds, the price is determined by the utility of the consumers who purchase the software but do not install updates. Such consumers are worse off from a higher α since it increases the probability of an attack by hackers. The reduced willingness to pay reduces the price and the firm's profits.

Part b(ii) of Proposition 2 considers a case in which before the increase in α Condition (C1) holds, but the new α is sufficiently high that (C1) ceases to hold. In such a case, the higher α induces a discontinuous increase in the optimal price and a discontinuous decrease in market share.

Part b(iii) and part c of Proposition 2 consider a situation in which condition (C1) does not hold, neither before nor after the increase in α. In this case, an increase in α has both costs and benefits

[26] Assumption A2 insures that there are such types.

to these consumers. The cost is that higher values of α increase the expected cost of installing updates. The benefit is that the expected damage goes down. 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 increases by $n(\gamma D - c)/2$.[27] Profits increase if and only if $\gamma > \hat{\gamma}$.

4. The Firm Does Not Disclose Vulnerabilities.

We now consider the case in which the firm does not disclose vulnerabilities and there is no regulation that requires it to do so. Are consumers necessarily worse off when the firm does not disclose vulnerabilities and does not issue updates? Clearly this depends on the type of consumer. High value consumers, who plan to install updates, will be worse off as they will be more vulnerable to hackers' activities. Such a policy increases the value of the product for low value consumers who would not install updates under disclosure of vulnerabilities. There will also be a group of “moderate-value” consumers who would install an update if issued, but are better off if the firm does not disclose vulnerabilities.

Since $W_{nd}(\theta)$ is increasing in θ (by Assumption A1), given the firm’s price, the consumers’ purchase decision can be characterized by a threshold type θ^{nd}, such that only consumers of type $\theta \geq \theta^{nd}$ will purchase the software.

Proposition 3 (No Disclosure): When the firm does not disclose vulnerabilities, the optimal price, market share, and profits are respectively $p^{nd} = T/2$, $1-\theta^{nd} = 1/2$, and $\pi^{nd} = T/4$.[28] ■

From Proposition 3, 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. Clearly, when the firm does not disclose vulnerabilities, changes in α or c have no effect on the equilibrium price or profits.

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

[28] Recall that $T \equiv v - \gamma Dn$.

5. 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.

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, $\theta^t=c/D\gamma$, such that consumers of type $\theta>\theta^t$ are better off when the firm discloses information, and consumers of type $\theta<\theta^t$ are better off when the firm does not disclose information. Consequently, there are two possible outcomes when firms can set their disclosure policy: The firm discloses vulnerabilities and sets a price such that $\theta>\theta^t$ and all consumers install updates. Alternatively, the firm sets a price such that $\theta<\theta^t$ and does not disclose vulnerabilities. Note that consumers of type $\theta\in[\hat{\theta},\theta^t]$ will install updates when available, but prefer a non-disclosure policy. (See Figure 4)

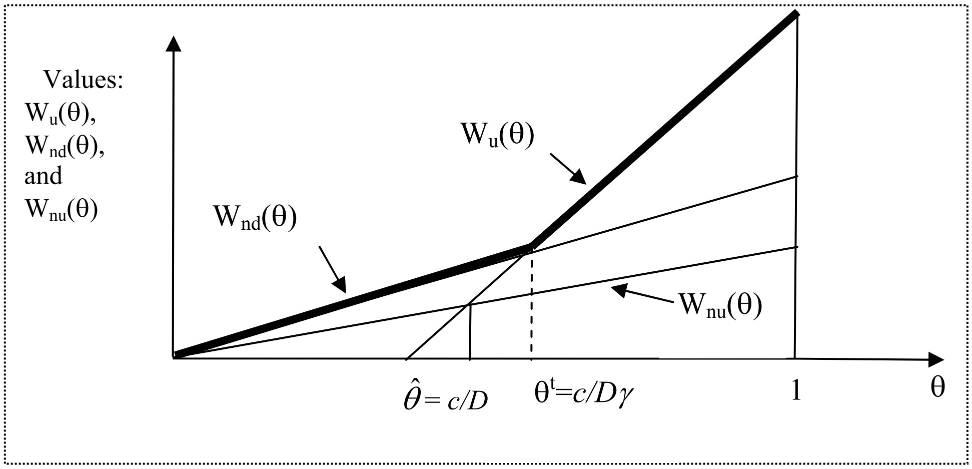

Figure 4: Willingness to pay under disclosure and non-disclosure

Proposition 4:

(i) The firm serves a larger market when it does not disclose vulnerabilities, but charges a higher price under disclosure of vulnerabilities.

(ii) The firm will choose not to disclose vulnerabilities if and only if $(1/2)(\alpha n(D\gamma - c)/2) < (\alpha nc/2Z)(Z - \alpha nc)/2$, which we will refer to as Condition (C2).[29]

(iii) Whenever $D<D^t \equiv [2-\alpha nc/(v-(1-\alpha)D^t n)]c$, the firm will not disclose vulnerabilities, regardless of the value of γ. Whenever $D>D^t$, there is a critical probability of hacker attack, $\gamma^t(n,\alpha,D,c,v)$, such that whenever $\gamma>\gamma^t$, the firm discloses vulnerabilities. ■

The left hand side of condition (C2) is the equilibrium market share when the firm does not disclose vulnerabilities multiplied by the difference in prices, while the right hand side is the difference in market shares multiplied by the equilibrium price when the firm discloses vulnerabilities. Hence, like Condition (C1), Condition (C2) says that the firm will not disclose vulnerabilities when the percentage gain in market share from doing so exceeds the percentage loss in price.

The intuition for part (iii) of the Proposition is as follows: A high value of γ means that attackers “learn” little from disclosure. We can view ($1-\gamma$) as a measure of what the hackers learn from disclosure. Without disclosure, hackers attack with probability γ, while following a disclosure of vulnerabilities, reverse engineering enables hackers to attack with probability one. In such cases, the firm will disclose vulnerabilities.[30] On the other hand, when is γ is relatively small, attackers learn a lot from the disclosure of vulnerabilities and the firm finds it optimal not to disclose vulnerabilities in this case.

Proposition 2 showed that profits can fall in α when the firm is required to disclose vulnerabilities. But when there is no mandatory disclosure policy, i.e., the firm can choose whether or not to disclose vulnerabilities, such an outcome is not possible. We discuss this issue further in section 7.2.

[29] Note that when Condition (C2) does not hold, Condition (C1) does not hold. Hence when the firm chooses disclosure, the equilibrium indeed involves all consumers installing updates. $\gamma>2c/D$ (or $D>2c/\gamma$) is a sufficient for Condition (C2) not to hold.

[30] This is true as long as $D>D^t$. Otherwise, the firm will not disclose vulnerabilities regardless of the value of γ.

6. Disclosure Policy, Regulation and Social Welfare

Some security experts recommend mandatory public disclosure of the discovery 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. As we discussed in the introduction, CERT/CC policy effectively mandates disclosure of vulnerabilities it reports to firms. In this section, we consider the effect of a regulatory policy requiring disclosure on social welfare.

We consider a regulator that can mandate the disclosure of vulnerabilities, but cannot set the price of the software or influence whether consumers install updates. Setting disclosure policy, however, does affect the market price as well as the market share. Since we assume no production costs, and since the price is a transfer from consumers to firms, social welfare is simply the integral of the willingness to pay for software.

When the firm discloses vulnerabilities and Condition (C1) holds, 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 updates. Summing up the surplus of these two groups of consumers gives us the Total Social Surplus, denoted SW_{nu}, in this case:

$$SW_{nu} = \int_{1/2}^{c/D} \mathrm{W}_{\mathrm{nu}}(\theta)d\theta + \int_{c/D}^{1} \mathrm{W}_{\mathrm{u}}(\theta)d\theta$$

$$= \int_{c/D}^{1} \{[v-(1-\alpha)\gamma Dn]\theta - \alpha nc\}d\theta + \int_{1/2}^{c/D} [v-(1-\alpha)\gamma Dn - \alpha Dn]\theta d\theta$$

$$= \frac{3Z}{8} - \alpha Dn\frac{(4c^2-D^2)}{8D^2} - \alpha nc\frac{(D-c)}{D}.$$

When the firm discloses vulnerabilities and Condition (C1) does not hold, the equilibrium is such that the firm sells only to consumers of type $\theta \in [\frac{(Z+\alpha nc)}{2Z}, 1]$ (See Proposition 2). Since these consumers also install updates, the total social surplus, denoted SW_u, is:

$$SW_u = \int_{1/2+\alpha nc/2Z}^{1} \mathrm{W_u}(\theta)d\theta = \int_{1/2+\alpha nc/2Z}^{1} \{[v \ -(1-\alpha)\gamma Dn]\theta - \alpha nc\}d\theta$$

$$= \frac{3Z}{8} + 3\frac{(\alpha nc)^2}{8Z} - 3\alpha nc/4 .$$

Finally, when the firms 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

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

Proposition 5:

(i) When Condition (C1) holds, the firm never discloses vulnerabilities while a regulator would mandate disclosure when the probability of attack is sufficiently large, i.e., for values of $\gamma > \gamma^{sI}$, where $\gamma^{sI} \equiv (8\beta - 4 - \beta^2)/3\beta^2$, and $1<\beta \equiv D/c<2$.[31]

(ii) When Condition (C1) does not hold, equilibrium disclosure policy is socially efficient. ■

Proposition 5 implies that equilibrium disclosure policy is not always socially optimal. Part (i) of the Proposition identifies circumstances under which the firm will choose not to disclose vulnerabilities, while welfare maximization requires such a disclosure. The intuition for (i) is that 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.

Proposition 5(ii) shows that equilibrium disclosure policy is welfare maximizing when Condition (C1) does not hold. In this case, 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 in this case. A regulator values an increased market share more than the firm does,

[31] For part (i) of Proposition 5, β must be between one and two because when $\beta>2$, condition (C1) does not hold.

because the firm obtains the full surplus only from the marginal consumer. In our setting, these opposing effects exactly cancel out. Hence, when Condition (C1) does not hold, the market outcome is efficient: A regulator would mandate disclosure whenever the firm would disclose vulnerabilities.[32]

Proposition 5 enables us to examine the effect of mandatory disclosure of vulnerabilities on welfare. When condition (C1) does not hold, the market outcome is efficient in our model. Hence, there are no benefits from mandatory disclosure. When Condition (C1) holds, mandatory disclosure improves welfare only when $\gamma > \gamma^{sI}$. However, mandatory disclosure would be welfare reducing when $\gamma < \gamma^{sI}$, since (in the absence of regulation) the firm does not disclose vulnerabilities and this is socially optimal.

Proposition 6 (Mandatory Disclosure):
(i) When Condition (C1) holds, mandatory disclosure decreases the equilibrium price.
(ii) When Condition (C1) does not hold, but Condition (C2) holds, mandatory disclosure increases the equilibrium price and reduces equilibrium number of consumers.
(iii) When Conditions (C1) and (C2) do not hold, mandatory disclosure has no effect on either the price or the number of consumers who purchase software. ■

The intuition for (i) is that when Condition (C1) holds, the firm would not disclose vulnerabilities in the absence of regulation. Since disclosure lowers the willingness to pay of all consumers in this case, it will lead to a lower equilibrium price. In case (ii), the firm would not disclose vulnerabilities in the absence of regulation. Since all consumers install updates under mandatory disclosure in this case, the firm serves a smaller market of higher quality consumers. Hence, mandatory disclosure leads to a higher equilibrium price and reduces the markets share. In case (iii), the firm indeed discloses vulnerabilities in the absence of regulation. Hence, mandatory disclosure has no effect in this case.

[32] Note that if, for example, θ was not uniformly distributed, the effects would not cancel out and the inefficiency (suboptimal or excess disclosure) would depend on the distribution of consumer types.

7. Ex-Ante and Ex-Post Investment

There are two types of investments the firm can undertake: (i) Investment that reduces the number of software vulnerabilities (i.e., investment to reduce n) and (ii) Investment that increases the probability that the firm will find problems before hackers (i.e., investment to increase 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 in quality.

7.1 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. This can be interpreted as an investment in reducing the number of software vulnerabilities before the software is sold. A reduction in n can be viewed as an increase in the quality of the product; thus it raises consumer willingness to pay for the software (See Figure 5).

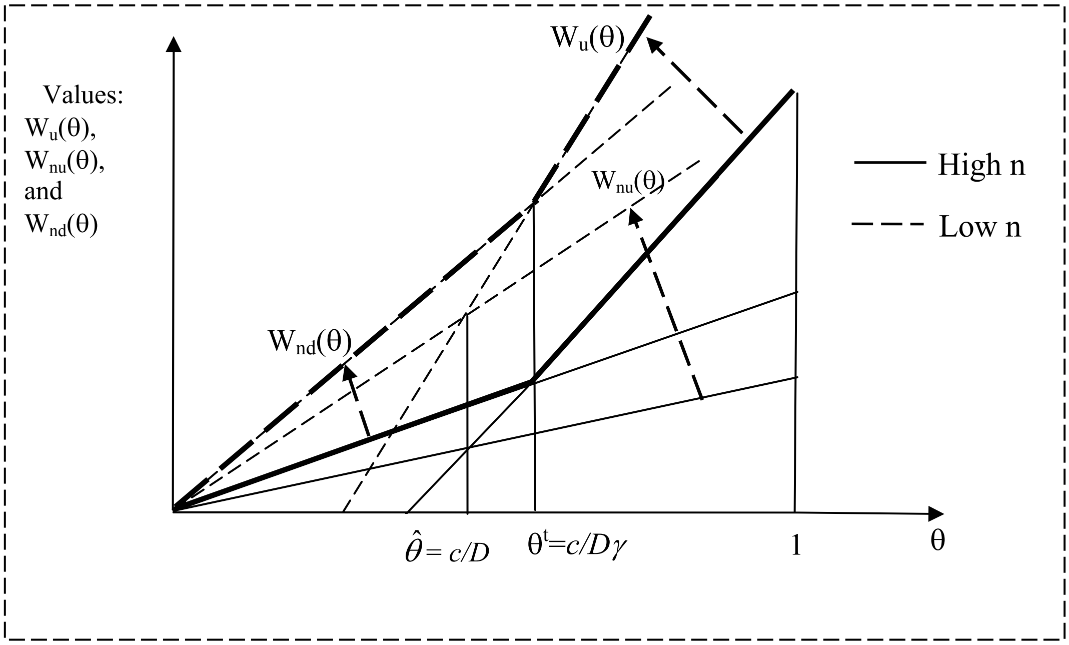

Figure 5: Effects of a decrease in n on $W_u(\theta)$, $W_{nu}(\theta)$,$W_{nd}(\theta)$

We now examine how a decline in the number of vulnerabilities (denoted Δn) affects equilibrium prices and disclosure policy.

Proposition 7 (Ex-Ante investment): Suppose that $D>D^t$. (i) When γ is sufficiently large, i.e., $\gamma>2c/D$, the firm discloses vulnerabilities. A reduction in n does not affect the disclosure policy of the firm. The reduction in n leads to an increase in equilibrium price, profits and consumers welfare. (ii) When γ is sufficiently small so that $\gamma<\gamma^t(n,\alpha)$,[33] the firm adopts a non-disclosure policy and a reduction in n will not change this policy. Lower n leads to an increase in equilibrium price, profits and consumers welfare. (iii) When $\gamma^t(n,\alpha)<\gamma<2c/D$ the firm discloses vulnerabilities. A small reduction in n (but where γ remains greater than $\gamma^t(n-\Delta n,\alpha)$ will not affect the disclosure policy but leads to an increase in equilibrium price, profits, and consumer welfare. On the other hand, a large decrease in the number of vulnerabilities (such that $\gamma<\gamma^t(n-\Delta n,\alpha)$) will induce a switch to non-disclosure and possibly to a lower equilibrium price. ■

Case (i) is intuitively derived from Proposition 4 because $\gamma>2c/D$ is a sufficient condition for the firm to disclose vulnerabilities. In case (ii) when γ is sufficiently small so that $\gamma<\gamma^t(n,\alpha)$, Proposition 4 shows that the firm finds it optimal not to disclose vulnerabilities. Furthermore since $\gamma^t(n,\alpha)$ is a decreasing function of n, a reduction in the number of software vulnerabilities implies a higher $\gamma^t(n,\alpha)$. Thus $\gamma<\gamma^t(n,\alpha)$ implies that $\gamma<\gamma^t(n-\Delta n,\alpha)$, which implies that regardless of the magnitude of the reduction in n, the firm finds it optimal to maintain its non-disclosure policy.

Finally, consider the intermediate case when $\gamma^t(n,\alpha)<\gamma<2c/D$. In this case the firm's optimal policy is to disclose vulnerabilities. But $\gamma^t(n,\alpha)$ is a decreasing function of n. Thus a reduction in n results in a higher $\gamma^t(n,\alpha)$ and, in particular, if the reduction is sufficiently large there will be a switch from case (iii) to case (ii) such that $\gamma^t(n-\Delta n,\alpha)>\gamma$. This induces the firm to change from a disclosure policy to a non disclosure policy.

Although a reduction in n is an improvement in software quality, our analysis indicates that the higher quality does not necessarily imply a higher equilibrium price in this case. If Δn is sufficiently large, there is a switch of the disclosure policy which, as Figure 5 indicates, may

[33] The parameters of interest in section 7 are α and n. Hence, we write $\gamma^t(n,\alpha)$ rather than $\gamma^t(n,\alpha,D,c,v)$.

result in either a higher or lower price. Formally, a reduction in n that leads to a regime change implies a lower price whenever $\Delta n < \alpha n(\gamma D-c)/\gamma D$, i.e., whenever Δn is not too large.

The intuition is that if the reduction in vulnerabilities is very large, the improvement in quality will offset the lower price associated with a "regime" change and the new equilibrium price will be higher. But when Δn is not too larger, i.e., $\Delta n < \alpha n(\gamma D-c)/\gamma D$, the regime change effect dominates the change in quality effect and the new equilibrium price is lower.

7.2 Ex-Post Investment: Increasing the Percent of Vulnerabilities the Firm Finds Before Hackers

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 this section, we investigate the effect of higher α on price, profits, and the firm's disclosure policy.

Proposition 8 (Ex-Post investment): Suppose that $D>D^t$. (i) When $\gamma>\gamma^t(n,\alpha)$, the firm would disclose vulnerabilities and an increase in α implies higher profits without any change in the firm's disclosure policy. (ii) When $\gamma<\gamma^t(n,\alpha)$, the firm does not disclose vulnerabilities. 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 and greater profits. ■

In case (i), the firm discloses vulnerabilities since $\gamma>\gamma^t(n,\alpha)$. Furthermore, there is no change in disclosure policy when α increases because $\partial\gamma^t(n,\alpha)/\partial\alpha<0$. Thus $\gamma>\gamma^t(n,\alpha)$ implies $\gamma>\gamma^t(n,\alpha+\Delta\alpha)$, which implies that disclosure is optimal regardless of the magnitude of the increase in α. Profits increase in α in case (i) because $\gamma^t(n,\alpha)> \hat{\gamma}(n,\alpha)$.

In case (ii), the optimal policy is not to disclose vulnerabilities since $\gamma<\gamma^t(n,\alpha)$. But since $\gamma^t(n,\alpha)$ is a decreasing function of α, an increase in α results in a lower $\gamma^t(n,\alpha)$. 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 to case (i). A switch from a non-disclosure policy to a disclosure policy takes place if $\gamma > \gamma^{t}(n, \alpha + \Delta\alpha)$.

8. 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 derived conditions under which a firm would disclose vulnerabilities. We then compared the equilibrium with the socially optimal disclosure policy and considered the effects of mandatory disclosure. Finally, we explored the effects of reductions in the number of vulnerabilities ex-ante and increases in the ex-post probability that the firm will find the vulnerabilities before hackers.

Our model sheds light on the effect of so-called “bug bounty” programs, in which firms offer rewards to users who identify and report vulnerabilities. These bounty programs have become quite popular and have attracted a lot of attention.[34] The effect of a bounty program can be interpreted in our setting as an increase in α.[35] Only users who install updates benefit from a bounty program.

From our analysis in section 7.2, the use of a bounty program has a positive effect on both profitability (and welfare) when γ is sufficiently large.[36] In such a case, the firm discloses vulnerabilities, the marginal consumer applies updates, and profits are increasing in α. When the firm chooses not to disclose vulnerabilities, the introduction of a bounty program has no effect unless it induces the firm to switch from non-disclosure to a disclosure policy.

[34] 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 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.”

[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] Since $SW_u = 3\pi^u/2$, an increase in profits increases Social Welfare as well.

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.

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*,

Meta Group Staff, “META Report: Security Vulnerability Disclosures,” January 2002, available at http://itmanagement.earthweb.com/it_res/article.php/947271

Samueslon, P., and S. Scotchmer, 2002, “The Law and the Economics of Reverse Engineering,” *Yale Law Journal*, 111: 1575-1663.

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:XX-XX.

Varian, H., 2004, “System Reliability and Free Riding,” available at http://www.ischool.berkeley.edu/~hal/Papers/2004/reliability.

Appendix: Proof of Propositions

Proof of Proposition 1:

(i) When $\theta^* < \hat{\theta}$, the marginal consumer does not install updates. In this case, profits are given by (4). Maximizing (4) yields $p^{nu} = S/2$, $1-\theta^{nu} = ½$, and $\pi^{nu} = S/4$.

(ii) When $\theta^* > \hat{\theta}$, the marginal consumer installs updates. In this case, profits are given by (5). Maximizing (5) yields $p^u = (Z - \alpha nc)/2$, $1-\theta^u = (Z - \alpha nc)/2Z$, and $\pi^u=(Z - \alpha nc)^2/4Z$.

(iii) Algebraic manipulation shows that $\pi^{nu} > \pi^u$ if and only if $\frac{1}{2}\frac{\alpha n(D-c)}{2} < \frac{\alpha nc}{2Z}\frac{(Z-\alpha nc)}{2}$. ■

Proof of Lemma 1:

(i) Condition (C1) can be rewritten $(D-c)/4c<(Z - \alpha nc)/4Z=(1-\theta^u)/2$.

Since $\partial(1-\theta^u)/\partial\alpha = -Tnc/2Z^2 <0$, the RHS decreases in α while the LHS does not depend on α. Hence, the RHS is minimized when $\alpha=1$. Hence Condition (C1) always holds if $D/c-1<1-nc/v$, since $Z=v$, when $\alpha=1$. Rewriting, Condition (C1) always holds when $D<(2-nc/v)c$.

(ii) Condition (C1) can be rewritten $\alpha n(D-2c) + \frac{(\alpha nc)^2}{Z}<0$. Since the second term is greater than zero, $D\geq 2c$ is a sufficient condition for Condition (C1) not to hold.

(iii) By continuity, when $(2-nc/v)c\leq D<2c$ there exists an α^c such that condition (C1) holds (does not hold) whenever $\alpha<(>)\alpha^c$. ■

Proof of Proposition 2:

a,b(i): $p^{nu} = [v - \gamma(1-\alpha)Dn- \alpha Dn]/2$. 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(ii),b(iii), c: $\partial p^u/\partial\alpha = \{\gamma Dn- nc \}/2 >0$, since $\gamma>c/D$ by assumption. $\partial(1-\theta^u)/\partial\alpha = -Tnc/2Z^2 <0$.

$\pi^u=\frac{(Z-\alpha nc)^2}{4Z}= ¼ \{Z - 2\alpha nc + \frac{(\alpha nc)^2}{Z}\}$. $\partial\pi^u/\partial\alpha=¼ \{\gamma Dn - 2nc + \alpha n^2c^2\frac{(2Z-\gamma\alpha Dn)}{Z^2}\}$.

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 $\gamma=\beta c/D$.

$\partial\pi^u/\partial\alpha=¼ \{\beta cn - 2nc + \alpha n^2c^2\frac{(2Z-\beta\alpha cn)}{Z^2}\}=\frac{cn}{4}\{\beta - 2 + \frac{\alpha nc}{Z}[2-\beta\frac{\alpha cn}{Z}]\}=$

$\frac{cn}{4}\{\beta - 2 + x(2-\beta x)\}$, where $x=\alpha nc/Z$. Since $Z>\alpha nc$, $x<1$. $\partial\pi^u/\partial\alpha>0 \Leftrightarrow$

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

Thus we have $\hat{\gamma}\equiv\frac{2c}{(1+\frac{\alpha nc}{Z})D}$ ($\hat{\gamma}$ is implicitly defined since Z is a function of γ), such that:

$\gamma>\hat{\gamma} \Rightarrow \partial\pi^u/\partial\alpha>0$

$\gamma=\hat{\gamma} \Rightarrow \partial\pi^u/\partial\alpha=0$

$\gamma<\hat{\gamma} \Rightarrow \partial\pi^u/\partial\alpha<0$ ■

Proof of Proposition 3:
Since the firm captures the surplus of the marginal consumer, the price and profits are as follows:

$$p^{nd}= [\theta^{np}v] -\gamma\theta^{np}Dn\equiv T\theta$$
$$\pi^{nd}(\theta^{nd}) = p^{nd}(\theta^{nd})\ (1-\theta^{nd})=\{\theta^{nd}v -\gamma\theta^{nd}Dn\ \}(1-\theta^{nd})\equiv T\theta(1-\theta)$$

Maximizing these profits yields the equations in the Proposition. ■

Proof of Proposition 4:
(i) $1-\theta^{u} = (Z - \alpha nc)/2Z < ½=1-\theta^{nd}$. The difference in prices is $p^{u}-p^{nd}=\dfrac{\alpha n(D\gamma-c)}{2}>0$ by A1.
(ii) Algebraic manipulation shows that $\pi^{nd} > \pi^{u}$ if and only if $\dfrac{1}{2}\dfrac{\alpha n(D\gamma-c)}{2}<\dfrac{\alpha nc}{2Z}\dfrac{(Z-\alpha nc)}{2}$.
(iii) The LHS of (C2) is negative when $\gamma=0$. Further, 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 $D<[2-\alpha nc/(v-(1-\alpha)Dn)]c\equiv D^{t}$. Thus, when $D<D^{t}$, the firm will not disclose vulnerabilities regardless of the value of γ. By continuity, whenever $D>D^{t}$, there exists a γ^{t} such that the firm is indifferent between disclosing and not disclosing vulnerabilities. When $\gamma<\gamma^{t}$, Condition (C2) holds and the firm will not disclose vulnerabilities. When $\gamma>\gamma^{t}$, Condition (C2) does not hold and the firm will disclose vulnerabilities. ■

Proof of Proposition 5:
(i) $SW_{nu}>SW_{nd}$ iff $\dfrac{3(v\ -\gamma Dn)}{8}+\dfrac{3\gamma\alpha Dn}{8}-\alpha Dn\{\dfrac{4c^2-D^2)}{8D^2}-\alpha nc\dfrac{(D-c)}{D}>\dfrac{3(v-\gamma Dn)}{8}$ or
$\gamma>\dfrac{8cD-4c^2-D^2}{3D^2}=\dfrac{8\beta-4-\beta^2}{3\beta^2}$ (where $D=\beta c$)
(ii) $SW_{nd}=\dfrac{3(v-\gamma Dn)}{8}=3\dfrac{T}{8}=\dfrac{3\pi^{nd}}{2}$

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

$$SW_{u}=\frac{3[v-\gamma(1-\alpha)Dn]}{8}-3\alpha nc/4+\frac{3(\alpha nc)^2}{8Z}=3\frac{Z}{8}-3\alpha nc/4+\frac{3(\alpha nc)^2}{8Z}=\frac{3\pi^{p}}{2}$$

Hence $SW_{u}> SW_{nd}$ iff $\pi^{u} > \pi^{nd}$. ■

Proof of Proposition 6:
(i) If C1 holds, 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 propositions 1 and 3, and since T>S).
$1-\theta^{nu}=1-\theta^{nd}=½$ (from propositions 1 and 3), and thus the equilibrium market share is unaffected.

(ii) If C1 does not hold and C2 holds, 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 propositions 1 and 3:
$p^u = (Z - \alpha nc)/2, p^{nd} = T/2 \Rightarrow p^u > p^{nd}$ since Z-αnc>T (from the definitions of Z, T and assumption A2).
$1-\theta^u = ½ - \alpha nc/2Z < ½ = 1-\theta^{nd}$, so the equilibrium market share decreases.

(iii) If C1 and C2 do not hold, mandatory disclosure does not cause a regime change. The equilibrium remains "disclosure" and all consumers install updates. Thus the price and market share remain p^u and $1-\theta^u$. ■

Proof of Proposition 7:

(i) γ>2c/D is a sufficient condition for the firm to disclose vulnerabilities (from the proof of Proposition 2).
(ii) From Proposition 4, the firm finds it optimal not to disclose vulnerabilities and $\gamma^t(n,\alpha)$ is a decreasing function of n (which we now show).

γ^t is implicitly defined by: $\gamma^t = \dfrac{2c}{D} - \dfrac{\alpha nc^2}{Z(\gamma^t)D}$, we denote $k \equiv \dfrac{\alpha c^2}{D}$, and thus

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

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

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

$$\Rightarrow \frac{\partial\gamma^t}{\partial n} < 0.$$

(iii) $\gamma > \gamma^t \Rightarrow$ The firm finds it optimal to disclose vulnerabilities, meaning $\pi^u > \pi^{nd} \Leftrightarrow$ Z(γD-c)>Zc-αnc² ⇔ -Z(2c-γD) + αnc²>0, where 2c-γD>0, since *γ<2c/D.* For a small decrease in n this inequality will continue to hold, because of continuity of the LHS. A small reduction in n results in a higher price since $\dfrac{\partial p^u}{\partial n} = -\gamma(1-\alpha)D - \alpha c < 0$. The result would also be higher profits for the firm (and higher consumer welfare) since $\gamma^t > \hat{\gamma}$.(See proof of Proposition 8)

It is obvious that as n decreases it is more likely that this inequality will cease to hold. Thus for a large decrease in n we will have $\pi^u < \pi^{nd}$, and the firm will switch to a non-disclosure regime.
For $\gamma < 2c/D$ there exists a critical value of n, denoted n^t, for which $\pi^u = \pi^{nd} \Leftrightarrow$

$$Z(\gamma D-2c)+\alpha nc^2=0 \Leftrightarrow n^t=\frac{v(2c-\gamma D)}{(1-\alpha)\gamma D(2c-\gamma D)+\alpha c^2}$$

$n>n^t \Rightarrow \pi^u > \pi^{nd}$ and the firm chooses disclosure.

$n< n^t \Rightarrow \pi^u < \pi^{nd}$ and the firm chooses non-disclosure.

Let n be initial value of the number of vulnerabilities, such that $n>n^t$, and the decline in the number of vulnerabilities is denoted Δn, such $n-\Delta n <n^t$.

The prices will be:

When $n>n^t$, $p^u=[Z(n)-\alpha nc]/2$

When $n-\Delta n <n^t$, $p^{nd}=[T(n-\Delta n)]/2$

$$p^{nd}< p^u \Leftrightarrow v-\gamma D(n-\Delta n)<v-\gamma(1-\alpha)Dn-\alpha nc \Leftrightarrow \Delta n<\frac{\alpha n(\gamma D-c)}{\gamma D} \qquad \blacksquare$$

Proof of Proposition 8:

(i) and (ii) follow from Proposition 2 and Proposition 4 because $\frac{\partial \gamma^t}{\partial \alpha}<0$ and $\gamma^t>\hat{\gamma}$.

Hence, we prove the Proposition by proving those two conditions:

We first show that $\gamma^t>\hat{\gamma}$: From the equation $\pi^u=\pi^{nd}$ we have γ^t implicitly defined: $\gamma^t=\frac{2c}{D}-\frac{\alpha nc^2}{Z(\gamma^t)D}$, and from $\frac{\partial \pi^u}{\partial \alpha}=0$ we have $\hat{\gamma}=\frac{2Z(\hat{\gamma})c}{(Z(\hat{\gamma})+\alpha nc)D}$.

Comparing the two, using simple algebra:

$$\frac{2c}{D}-\frac{\alpha nc^2}{Z(\gamma^t)D}\overset{?}{>}\frac{2Z(\hat{\gamma})c}{(Z(\hat{\gamma})+\alpha nc)D}$$

$$2Z(\gamma^t)Z(\hat{\gamma})+2Z(\gamma^t)\alpha nc-\alpha ncZ(\hat{\gamma})-(\alpha nc)^2\overset{?}{>}2Z(\gamma^t)Z(\hat{\gamma})$$

$$2Z(\gamma^t)-Z(\hat{\gamma})-\alpha nc\overset{?}{>}0$$

$$Z(2\gamma^t-\hat{\gamma})-\alpha nc\overset{?}{>}0$$

From assumption A1 we have $\forall 0<\gamma<1: Z-\alpha nc>0$.

If 0< $2\gamma^t-\hat{\gamma}$ <1 then the above inequality holds and we have $\gamma^t>\hat{\gamma}$. Otherwise there are two possibilities:

(i) If $2\gamma^t-\hat{\gamma}$ >1, then $\gamma^t>0.5(\hat{\gamma}+1)$. If $\gamma^t<\hat{\gamma}\Rightarrow\gamma^t>1$. This contradicts our assumption that $0<\gamma<1$. Thus, in this case, $\gamma^t>\hat{\gamma}$.

(ii) If $2\gamma^t-\hat{\gamma}$ <0, then $\hat{\gamma}>2\gamma^t>\gamma^t$. But, Z($2\gamma^t-\hat{\gamma}$) will be greater than Z(γ) when γ>0. Thus the above inequality [$Z(2\gamma^t-\hat{\gamma})-\alpha nc\overset{?}{>}0$] will continue to hold, meaning $\gamma^t>\hat{\gamma}$. But this is a contradiction, since $2\gamma^t-\hat{\gamma}$ <0 implies $\gamma^t<\hat{\gamma}$.

Thus, the only case is 0< $2\gamma^t-\hat{\gamma}$ <1, and $\gamma^t>\hat{\gamma}$.

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

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

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

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

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

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

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

■