The Countervailing Incentive of Restricted Patch Distribution: Economic and Policy Implications

Mohammad S. Rahman Karthik Kannan
Mohit Tawarmalani
Krannert School of Management, Purdue University, West Lafayette, Indiana 47907
{mrahman, kkarthik, mtawarma}@mgmt.purdue.edu

March 2007

Abstract

Traditionally, the government has been the sole entity to enforce anti-piracy measures. Of late, software vendors are attempting to thwart piracy of their products by providing patches only to legal users. By doing so, a vendor can vertically differentiate the legal copy from the pirated copy. It is not clear if the vendor's differentiation strategy complements or substitutes the government's effort with respect to social welfare. We study this issue by analyzing the impact of a monopolistic vendor's action to restrict patches on both the vendor's profit and the social welfare. Two key distinguishing features of our model are: (i) we endogenize the hacker activity and, therefore, the loss suffered by the users, and (ii) we also endogenize the quality of the patch developed by the vendor. Based on our analysis, we find that a monopolist does not always benefit from vertical differentiation. More specifically, when the government's anti-piracy effort is intense and the cost of developing a good quality patch is high, the vendor does not benefit from vertical differentiation. Another interesting result of our analysis is that, by strategically utilizing the hacker's activity, it is possible to improve social welfare relative to that when the patch is universally distributed.

Key words : information security; patch distribution; countervailing incentive; public policy

1 Introduction

Software piracy and unpatched software are important problems for policy makers, businesses, and users. Typically, the government has been the sole entity to enforce anti-piracy measures. Lately, software vendors have tried to mitigate piracy of their products by requiring authentication of the software before offering security patches (Png et al., 2006; August and Tunca, 2006a). One prominent example of such practice is Windows Genuine Advantage. Under this scheme, Microsoft checks for legitimate copies of Windows before any updates or patches are downloaded.[1] Clearly, this practice vertically differentiates a legitimate copy from a pirated copy. It has been noted that restricting patches through Windows Genuine Advantage has led to an increase in the number of systems that have been attacked (Naraine, 2005). Given that security patches play a crucial role in de-motivating the hackers,[2] the interplay between piracy and information security becomes evident. A hacker's motivation to exert effort increases as the number of users increases, especially as more vulnerable systems are present. This is because the marginal benefit from exerting effort increases with the number of systems that can be exploited.[3] Since piracy leads to a larger user base, software piracy, perhaps implicitly, has important implications for information security.

The conundrum is whether or not a software vendor should make security patches available to pirates. On the one hand, if pirated systems are not patched, the overall hacking activity increases which has adverse implications for legal users (assuming patches are not perfect). The adverse implication, in turn, leaves lesser surplus (for the legal users) that the vendor can extract. On the other hand, if pirated copies are also supported with security patches, the demand for legal copies is cannibalized. The user's decision to pirate or not also depends on the government's anti-piracy measures. Obviously, the demand for the legal copy increases with the intensity of the anti-piracy effort. For a government (social planner), which seeks to maximize the welfare of the users–both

[1] See http://www.microsoft.com/genuine/downloads/whyValidate.aspx for more details.

[2] In this paper, we focus on unethical hackers who attack users with malicious intents. For brevity, we simply refer to them as "hackers".

[3] Once the hacker learns about a specific vulnerability, he can attack all other systems with that vulnerability at practically no cost.

legal and illegal–as well as the vendor, a related question is: how does the intensity of the anti-piracy effort vary with the vendor's decision to restrict the patch?

It is not ex ante clear if it is welfare improving to complement the vendor's action with anti-piracy effort. If indeed the restricted patch distribution is a substitute of the social planner's action (i.e., improves social welfare), then the government does not need to exert any anti-piracy effort. However, if the welfare decreases when the patch is restricted,, then the social planner should resort to appropriate regulations that achieves the desired objective. Thus, from a social planner's perspective, understanding the implications of a vendor's decision to restrict patches only to legal users is crucial. In order to provide insights to the social planner, we develop and analyze a game-theoretic model. Two key distinguishing features of our model are: (i) we endogenize the hacker activity and, therefore, the loss suffered by the users, and (ii) we also endogenize the quality of the patch developed by the vendor.

Our analysis identifies two different effects of the hacker's activity. The adverse effect, which occurs independent of the vendor's decision to restrict the patch, decreases the welfare of the legal users that the monopolist can extract. In contrast, the countervailing effect occurs only when the patches are restricted. The significant negative impact of the hacker's activity on the pirates compare to that on the legal users helps the vendor vertically differentiate the legal copy from the illegal one. We find it surprising that the hacker's activity, which destroys consumer welfare, in fact improves the social welfare, when the anti-piracy effort is low. This is so because the decision to restrict the patch incentivizes the monopolist to develop a better quality patch that leads a higher surplus for the legal users. Correspondingly, the countervailing effect dominates the adverse effect.

Another surprising result is that, although one would expect a monopolist to always benefit from the vertical differentiation, it is not so here. When the cost of developing a quality patch is high, the monopolist does not benefit from the vertical differentiation. This is because, in such a case, the vendor chooses a low quality patch which decreases the differentiation. This, in turn, results in the adverse effect dominating the countervailing effect. Also, it is counterintuitive to note that the vendor does not have incentive to complement the intense anti-piracy action with the decision to restrict the patch. In this case, complementing the government's anti-piracy effort with the vendor's decision to restrict reduces the legal user's relative willingness to pay compared to the case where only the government's anti-piracy action is utilized.

This paper is organized as follows. In §2, we review the extant literature most relevant to this topic. Following that, in §3, we describe our model. In §4, we present our equilibrium analysis and compare the two policies - providing a universal patch as opposed to restricting the patch only to the legal users. Social welfare implications are discussed in §5. Finally, we present our concluding remarks in §7.

2 Literature Review

This paper overlaps two different research streams, information security and piracy. Information security is not only a technical problem but also an issue of economic incentives (Anderson, 2001). There are a number of papers that focus on the economic aspects of information security. Gordon et al. (2003) demonstrate how incentive issues surrounding information sharing in Information Sharing & Analysis Centers (ISACs) are similar in spirit to those in trade associations. They highlight the impact of information sharing on security investment and information security. They also provide insights regarding free-riding, which potentially poses serious challenges for information sharing. Relatedly, Gal-Or and Ghose (2005) focus on the competitive implications of sharing information about security breaches and security investments. Their results highlight how information sharing complements security investment. Gordon and Loeb (2002) analyze an economic model of information security investment. Their analytical results contend that the optimal level of information security spending does not always increase with the expected loss from attacks. Also, the level of security spending needs to be a small fraction of the expected loss from attacks. Cavusoglu et al. (2005) discuss the value of implementing Intrusion Detection Systems within firms.

Many papers have focused on analyzing different incentives involved in discovering, disclosing, and patching vulnerabilities. For instance, Arora et al. (2004) provide an economic decision-making framework for disclosing vulnerabilities. They show that vulnerability disclosures expedite the response from large vendors and subsequently benefit software users. Also, Kannan and Telang (2005) demonstrate that a passive CERT-type mechanism almost always generates better social outcome in comparison to a market-based mechanism for vulnerability disclosure. Arora et al. (2006) show that when a market is big, a producer is better-off releasing a buggier software early and patching it later. The researchers suggest that, in the presence of competition, a vendor offering high value to customers is better off releasing a buggier product early. This stream of research suggests that most of the software released are vulnerable and need patching for appropriate security. As a result, any policy regarding security patch distribution has significant implications for users, vendor, and welfare.

Png et al. (2006) consider the strategic interaction between end-users in taking security precautions, and interaction between end-users and hackers. But , their work does not focus on the economic and policy implications of piracy on the interaction among end-users, hacker, and vendor. Recently, August and Tunca (2006b) consider users' incentive to patch security flaws. They find that subsidy based patching policy performs better than mandatory or tax based patching policy. They contend that the more users patch the system the better it is for the overall network security. They suggest that by making patching cost low (by making it easy for users to patch and providing reliable patch), a vendor or social planner can improve network security.

Most prior work on software piracy analyze the impact of piracy on legitimate producer's sales and profit. A common consensus is that a producer may have the incentive not to eliminate piracy from the market (Chen and Png, 2003; Gopal and Sanders, 1997; Shy, 2001). Piracy generates network externality benefits which lead to increased demand for legitimate version (Conner and Rumelt, 1991; Shy, 2001). In addition, a monopolist can commit not to decrease price in future and enjoy increased profit (Takeyama, 1994). In the similar vein, Varian (2000) argues that sharing or copying information goods can lead to increased profit for a producer if the transaction cost of sharing is lower than the marginal cost of production. Further, he argues that when sharing paves the way to distinguish between high-value and low-value customers, a producer's profit increases. The stream of research on piracy has also examined social welfare implications, and the results are typically inconclusive. Generally, strict rules to combat piracy increase the producer's profit while reducing the benefits of utilizing already developed products (Chen and Png, 2003). Chen and Png (2003) contend that from the social welfare perspective, it is better to manage piracy through price cuts than strict enforcement. The natural question that arises is what the is welfare implication of managing piracy through restricted patch distribution. The current paper attempts to address this question.

We recently encountered an independent work by August and Tunca (2006a) that also considers the implications of restricting patch distribution. They show that a vendor benefits from restricting patch distribution to only legal users if the software is highly risky and anti-piracy actions are mild, or the population's tendency to pirate is high. They also discuss the social welfare implications of restricted patch distribution. While our focus is similar to theirs, our model set-up is not. Unlike their work, we do endogenize both the hacker activity and the quality choice of the patch. Our motivation to endogenize the hacker activity is based on the anecdotal evidence which shows that a software with a large user base tends to attract more hacker activity.[4] As a consequence of endogenizing these variables, the problem becomes more involved and many aspects of the results are different. Thus, we provide insights into a framework where a hacker, a vendor, and users strategically interact with each others.

3 Model

Our model involves four participants, an anti-piracy agency (also referred to as the *government*), a software vendor, a hacker, and software users. We investigate the problem in the context of a monopolistic software vendor. The sequence of moves in our formulation is as follows. First, the government chooses the anti-piracy effort level. Then, the vendor decides on both the price and the

[4] For instance, a rational hacker is predisposed to attacking Windows users more often than Apple users. This is because, ceteris paribus, the expected payoff is higher in attacking windows due to the larger user base.

patch quality, as well as whether or not to make the patch available to the pirates. Subsequently, the hacker exerts an effort level to attack the systems and the users decide to pirate or purchase the software.

We denote the effort exerted by the government by $\alpha \in [0,1]$. Here, $\alpha = 0$ implies that the government does not exert any anti-piracy effort, whereas $\alpha = 1$ represents complete piracy elimination by the government. Let p be the price for the software and $x \in [0,1]$ be the patch quality decided by the vendor. In our model, $x = 1$ denotes that patch is of the highest quality and able to deter the hacker's attack with certainty. We assume that the vendor always provides the patch to users who have purchased the software. His decision to limit the access to the patch is represented by $z \in \{0,1\}$. Here, $z = 1$ denotes that the vendor makes the same patch information (including the patch itself) available to all users, even to the pirating users. On the other hand, $z = 0$ denotes that the patch is only available to legal users. For the hacker, the effort exerted to attack the systems is the decision variable. Specifically, we let that variable to be β or the probability of finding a vulnerability. The users, as mentioned earlier, have the option to pirate or buy the product.

We assume that software users are heterogeneous in terms of the intrinsic value they derive from the software. We let the user type, $\theta \in [0,1]$, be distributed according to a distribution $F(\theta)$. Normalizing the total number of software users in the market to one is without the loss of generality. Similar to that in Kannan and Telang (2005), we assume that the intrinsic value of the software is θ_i^2 for a software user of type θ_i. This value diminishes as the hacker gains access to her machine. If β is the effort exerted by the hacker and x is the quality of the patch, then $\beta(1-x)$ represents the probability with which a machine is compromised due to the hacker's effort. The expected consumer surplus for type θ_i from *buying* the legal software is

$$CS_b(\theta_i) = (1 - \beta(1-x))\,\theta_i^2 - p. \tag{1}$$

Note that the legal user's expected loss incurred from a successful breach depends on the hacker's effort and the patch quality since we assume that the vendor always makes the patch available to the legal users. In the above expression, the following condition will be required: $\frac{\partial CS_b(\theta_i)}{\partial \beta} \leq 0$. It implies that as the effort exerted by the hacker goes up, the consumer surplus decreases. Also, we require that $\frac{\partial CS_b(\theta_i)}{\partial x} \geq 0$. This implies that as the patch quality improves, the consumer surplus increases.

We let the pirated product to be an inferior but vertically differentiated substitute for the legal version. The vendor achieves the vertical differentiation by controlling the patch availability z. As a consequence of this control, the probability with which the machine is compromised increases. We represent the probability that a pirated software is compromised by $\beta(1 - xz)$. In this expression, when $z = 1$, no vertical differentiation is achieved by the vendor whereas $z = 0$ makes the pirated product inferior. Additionally, the government's anti-piracy effort also decreases the utility by a factor of $(1 - \alpha)$. It can be interpreted as the probability with which the pirated user may be subject to legal actions. Assuming that the cost for the pirated version is zero, the consumer surplus for type θ_i from *pirating* is:

$$CS_p(\theta_i) = (1 - \beta(1 - x\ z))\ \theta_i^2\ (1 - \alpha). \tag{2}$$

In this expression, when $z = 0$, the probability with which the hacker gains control of the machine is equal to β . In other words, if the vendor controls the patch availability more tightly, the utility for the pirated copy decreases. Note that when the vendor makes the patch available to everybody $z = 1$ and when the government does not exert any anti-piracy effort $\alpha = 0$, utilities from both the legal and the pirated versions of the product are identical. Also, notice that α serves to vertically differentiate the legal version from the pirated version independent of the value of z, whereas β serves to vertically differentiate when $z = 0$.

Let $\bar{\theta}$ be the user type who is indifferent between pirating and buying the software. By equating

(1) and (2), the indifference type computed to be

$$\bar{\theta} = \frac{\sqrt{p}}{\sqrt{\alpha - \alpha\,\beta + \beta\,x - (1-\alpha)\;\beta\,x\,z}}. \tag{3}$$

Type $\theta_i > \bar{\theta}$ will buy the software and the others pirate.

3.1 Hacker's Profit

In this section, we characterize the hacker's expected profit function. Similar to Kannan and Telang (2005), we assume that the hacker's gain from attacking is less than the loss incurred by the user. Let the hacker gain θ_i from successfully breaking into the system of user-type θ_i. The success of breaking-in is different between the pirated users and the legal users. Recall that with probability $(1-x)$ the legal users are protected from hacker's effort β. Subsequently, the probability of a successful break-in is simply $\beta(1-x)$ for the legal users. On the other hand, the probability of breaking into the pirated versions of the machines increases if the patch availability is restricted to the legal users. Specifically, $\beta(1-x\;z)$ is the probability that the hacker breaks into the pirated machine when he exerts an effort of β. If the hacker's cost is $C(\beta)$, the hacker's objective function is $\max_\beta \Pi_h$, where Π_h, the expected profit for the hacker is:

$$\Pi_h = \beta(1-x\;z)\int_0^{\bar{\theta}} \theta\;\mathrm{d}\theta + \beta(1-x)\int_{\bar{\theta}}^1 \theta\;\mathrm{d}\theta - C(\beta). \tag{4}$$

Note that it is extremely costly to exert effort with which a system will be compromised with certainty. In contrast, when the hacker exerts no effort, he does not incur any cost. In our model, we use the commonly used logarithmic cost function, $C(\beta) = -M\log(1-\beta)$, for the hacker to exert effort and attain the probability of success β. As a result of this assumption, the cost of $\beta = 1$ is infinity. In the above, M is the cost of exerting effort. Substituting for $C(\beta)$ in equation (4) and integrating by parts, we obtain

$$\Pi_h \;=\; \beta\,(1-x\,z)\big(\bar{\theta}-\int_0^{\bar{\theta}} F(\theta)\,\mathrm{d}\theta\big)+\beta(1-x)\,\big((1-\bar{\theta}F(\bar{\theta}))-\int_{\bar{\theta}}^{1} F(\theta)\,\mathrm{d}\theta\big)+M\log(1-\beta). \quad (5)$$

3.2 Vendor’s Profit Function

Since we have normalized the total number of users to one, the demand that the vendor encounters for its software, η, is

$$\eta = 1 - F(\bar{\theta}). \quad (6)$$

The vendor is assumed to incur a negligible marginal cost to produce the software. However, the vendor incurs a cost $K(x)$ in order to improve the patch quality and decrease the effective vulnerability of a patched system. As a result, the software vendor maximizes the following profit function:

$$\max_{p,x,z} \quad \eta p - K(x). \quad (7)$$

The term ηp corresponds to the revenue that the vendor generates from selling the product to legal users at price p. In addition to maximizing the price and patch quality, the vendor decides whether or not to restrict access to the patch to only legal users. Similar to hacker’s cost function, we use logarithmic cost function $K(x) = -L\ \log(1-x)$ for the vendor. The vendor incurs L for achieving patch quality x.

4 Equilibrium Analysis

We first sketch the behavior of software users, the hacker, and the vendor assuming the government chooses an exogenous level of anti-piracy effort α. Following the government’s action in period 0, the vendor sets the (p, x) pair as well as chooses z in the first period, and both the hacker and users react in the second period. An appropriate equilibrium concept for such games is Subgame Perfect Nash Equilibrium (SPNE) (Fudenberg and Tirole, 1991). We use backward induction in solving the game. Therefore, we first solve for the reactions of the hacker and the software users for a given (p, x, z) triplet. Next, we solve for the optimal price, patch quality, and patch distribution decision. Finally, we calculate the welfare-metrics for each scenarios.

Let us now characterize the hacker's optimal action, β^*. The optimal hacker effort, β^*, is a solution to the implicit equation (5) that requires some functional form assumption for $F(\theta)$. For the sake of simplicity, we assume θ be distributed uniformly on the interval $[0, 1]$. This implies that $F(\theta) = \theta$. Substituting for $F(\theta)$ and simplifying the equation, we obtain

$$\begin{aligned} \Pi_h &= \beta\,(1 - x\,z)\big(\bar{\theta} - \int_0^{\bar{\theta}} \theta\,\mathrm{d}\theta\big) + \beta(1-x)\,\big((1-\bar{\theta}^2) - \int_{\bar{\theta}}^1 \theta\,\mathrm{d}\theta\big) + M\log(1-\beta) \\ &= \frac{1}{2}\beta\left(1 + x\left(\bar{\theta}^2 - z\bar{\theta}^2 - 1\right)\right) + M\log(1-\beta). \end{aligned} \tag{8}$$

To obtain the optimal effort level of the hacker (β^*), we take the first-order condition on hacker's expected profit expression (8), substitute $\bar{\theta}$ from (3), and solve the resulting equation and simplify:

$$\begin{aligned} \beta^* = & \\ & \frac{2\,\alpha\,(M + x - 1) + (1 - 2\,M - p - x)\,x\,(1-z) + \alpha\,(1 - 2\,M - x)\,x\,z}{2\,(1-x)\,(x - \alpha - (1-\alpha)\,x\,z)} \\ + & \frac{\sqrt{4\,(1-x)\,(\alpha\,(1 - 2\,M - x) + p\,x\,(1-z))\,(x - \alpha - (1-\alpha)\,x\,z) + (2\,\alpha\,(M + x - 1) + (1 - 2\,M - p - x)\,x\,(1-z) - \alpha\,x\,(2\,M + x - 1)\,z)^2}}{2\,(1-x)\,(x - \alpha - (1-\alpha)\,x\,z)}. \end{aligned}$$

Subsequently, substituting the optimal hacker effort β^* in (3), we can compute the $\bar{\theta^*}$:

$$\begin{aligned} \bar{\theta^*} = & \\ & \frac{\sqrt{2}\,\sqrt{p}}{\sqrt{\frac{2\,\alpha\,M + (1-2\,M-p-x)\,x\,(1-z) + \alpha\,(1-2\,M-x)\,x\,z + \sqrt{4\,(1-x)\,(\alpha\,(1-2\,M-x) + p\,x\,(1-z))\,(x-\alpha-(1-\alpha)\,x\,z) + (2\,\alpha\,(M+x-1) + (1-2\,M-p-x)\,x\,(1-z) + \alpha\,(1-2\,M-x)\,x\,z)^2}}{1-x}}}. \end{aligned}$$

We can conclude the following from the optimal hacker strategy $\forall\ M \in [\frac{1}{2}, \infty]$, $\beta^* \leq 0$.

4.1 Patch is available to all users

The vendor maximizes $\eta\,p + L\ \log(1-x)$ subject to the constraint that $\beta^* \geq 0$. Note that when we maximize vendor's profit expression (7), we need to ensure that the hacker exerts non-negative effort in the equilibrium. By substituting $\bar{\theta}^{*}$, and setting $z = 1$ in these expressions, we obtain vendor's decision problem when the patch is available to *everyone*.

The corresponding profit of the vendor when $z = 1$ is:

$$\Pi_v^{\text{z=1}} = \begin{cases} \dfrac{8M\alpha}{27} + \dfrac{4\alpha}{27} - L + L\log\left(\dfrac{27L}{4\alpha}\right) & \beta^*_{\text{z=1}} > 0 \\ \dfrac{4\alpha}{27} + L\log(2M) & \beta^*_{\text{z=1}} = 0 \end{cases}$$

We observe the following properties in this ($p^*_{\text{z=1}}$, $x^*_{\text{z=1}}$) pair:

- Both p and x are increasing in α, the government's anti-piracy measure (i.e., $\frac{\partial p^*_{\text{z=1}}}{\partial \alpha} > 0$ and $\frac{\partial x^*_{\text{z=1}}}{\partial \alpha} \geq 0$).

- When the hacker exerts effort, as L increases (i.e., the cost of patch quality increases), both p and x decreases (i.e., $\frac{\partial p^*_{\text{z=1}}}{\partial L} < 0$ and $\frac{\partial x^*_{\text{z=1}}}{\partial L} < 0$).

- The effect of M on x varies depending on whether the constraint on β is binding or not. If the hacker exerts effort, the patch quality, x, provided by the vendor is independent of the cost incurred by the hacker. It is so because the patch is available to everyone (legal users and pirates). Hence, when the vendor influences the hacker's effort by altering x, it affects both legal users and pirates equally; thus it fails to create a strategic advantage. Notice that, *ceteris paribus*, as x increases, the demand for the legal software increases. However, if x increases, $\beta^*_{\text{z=1}}$ decreases, which, in turn, reduces the demand for legal software. Consequently, there is no benefit gained from increasing x. Hence, the vendor does not change x with M.

- When $\beta = 0$, as M increases, the value of x decreases. In this case, the vendor only needs to maintain a patch quality that is sufficient to keep the hacker out the market.

Now, by substituting ($p^*_{\text{z=1}}$,$x^*_{\text{z=1}}$) pair and setting $z = 1$, we find the equilibrium hacker effort:

$$\beta^*_{\text{z=1}} = \begin{cases} 1 - \dfrac{8\alpha M}{27L} & \lambda^*_{\text{z=1}} \leq 0 \\ 0 & \text{otherwise} \end{cases}$$

We observe the following properties in $\beta^*_{\text{Z=1}}$:

- The hacker's effort, $\beta^*_{\text{Z=1}}$, is decreasing in M (i.e., the cost of exerting hacking effort).
- As α increases (i.e., the anti-piracy effort intensifies), the optimal hacking effort level decreases. This is because with increasing α, the vendor has a higher incentive to provide a better quality patch, thus makes hacking more difficult.
- Finally, as L increases (i.e., the cost of patch quality increases), $\beta^*_{\text{Z=1}}$ increases.

Note that by the construction of the model $x \in [0, 1]$. The following can be derived immediately:$\forall \quad L \in [0, \frac{4\alpha}{27}]$, $0 \leq x \leq 1$.

4.2 Patch is available to only legal users

In this section, we analyze the case where the vendor restricts the access to the patch to only *legal users*. In optimizing the profit expression (7), in the equilibrium, we once again need to ensure that the hacker exerts non-negative effort. By substituting $\bar{\theta}^*$, and setting $z = 0$ in (7) and (8), we get:

$$\begin{aligned} &\max_{p,x} \quad p\left(1-\frac{\sqrt{2}\sqrt{p}}{\sqrt{\frac{x^2+2Mx+px-x-2\alpha M+\sqrt{4\alpha^2M^2-4\alpha x(2M+p+x-1)M+x^2\left(4M^2+4(p+x-1)M+(p-x+1)^2\right)}}{x-1}}}\right)+L\log(1-x) \\ &\text{s. t.} \quad \frac{2\alpha(1-M-x)-x+x(2M+p+x)-\sqrt{4\alpha^2M^2-4\alpha x(2M+p+x-1)M+x^2\left(4M^2+4(p+x-1)M+(p-x+1)^2\right)}}{2(1-x)(\alpha-x)} \geq 0 \end{aligned}$$

This optimization problem is analytically intractable.[5] As a result, we use numerical analysis to characterize the properties of the optimal (p, x) pair and to subsequently provide intuitions. Note that we have a compact solution space, which enables us to search the entire solution space and determine the optimal (p, x) for any given triplet $\{\alpha, M, L\}$.

In the numerical analysis, both α and M were initialized to 0.01, and each were incremented in steps of 0.01. Recall that the maximum value of L is limited to $\frac{4\alpha}{27}$. In order to study the effect of L

[5] We took first-order condition of the objective function with respect to p and x, ignoring the constraint. Even without considering the constraint, we found that the degree of the polynomial is $\frac{5}{2}$, for which no known direct factorization technique exists.

for low values of α, we choose to increment L in steps of 0.001. Finally, p and x were initialized to be 0.01, and were incremented by 0.01 at each step. For a given $\{\alpha,\, M,\, L\}$, we determine the optimal (p, x) pair that maximizes the vendor's profit requiring β^* to be non-negative. If β^* was non-positive, we computed the vendor's profit by substituting $\beta = 0$ in (3) and using this value in (7). Thus, for each combination of $\{\alpha, M,$ and $L\}$, we computed $100 * 100 = 10,000$ values of vendor profit and picked the maximum profit.[6] The (p, x) pair that leads to this maximum profit is the approximate equilibrium price and patch quality for the respective values of $\{\alpha\ ,\, M,\, L\}$.

Notice that this algorithm is not dependent on the value of z. So, we validated the algorithm by verifying the outcome of our algorithm with the analytical results when $z = 1$. Then we repeated the algorithm for $z = 0$.

We observe from our numerical analysis that sensitivity of $p^*_{\text{z=0}}$,$x^*_{\text{z=0}}$, and $\beta^*_{\text{z=0}}$ with respect to α and L are directionally similar in nature to the case when $z = 1$. More specifically, as the anti-piracy effort intensifies, the software vendor has the incentive to increase the price as well as the patch quality. This is intuitive, because intense anti-piracy effort forces users to obtain a legal version and allows the vendor to strategically benefit from the differentiation between the legal and the pirated version. Also, as more users switch to legal version due to anti-piracy effort, the hacker generates less benefit. This is because more users are protected with the patch. If the cost of patch quality, L, increases, the patch quality declines; consequently, the hacker benefits from increases in L.

Before we analyze the role of M, it is important to understand the effects of β. As such, β decreases the consumer surplus which the vendor can potentially extract. We call this effect the "adverse effect" of hacking. However, β also serves to vertically differentiate the legal version from the pirated version and aids the vendor. Recall that without the patch, the effect of β is more on $CS_p(\theta_i)$ than $CS_b(\theta_i)$. We call this effect the "countervailing effect"of hacking. As the hacker's cost M increases, as expected of β, the hacker exerts less effort. Interestingly, the effect of M on

[6] We used Matlab 7.1 to determine the optimal price and quality pair as well as other values. We could not search the solution space at further granular level due to the time requirement of the algorithm. Notwithstanding, we do not expect any changes in the qualitative nature of the intuition gained from the current analysis.

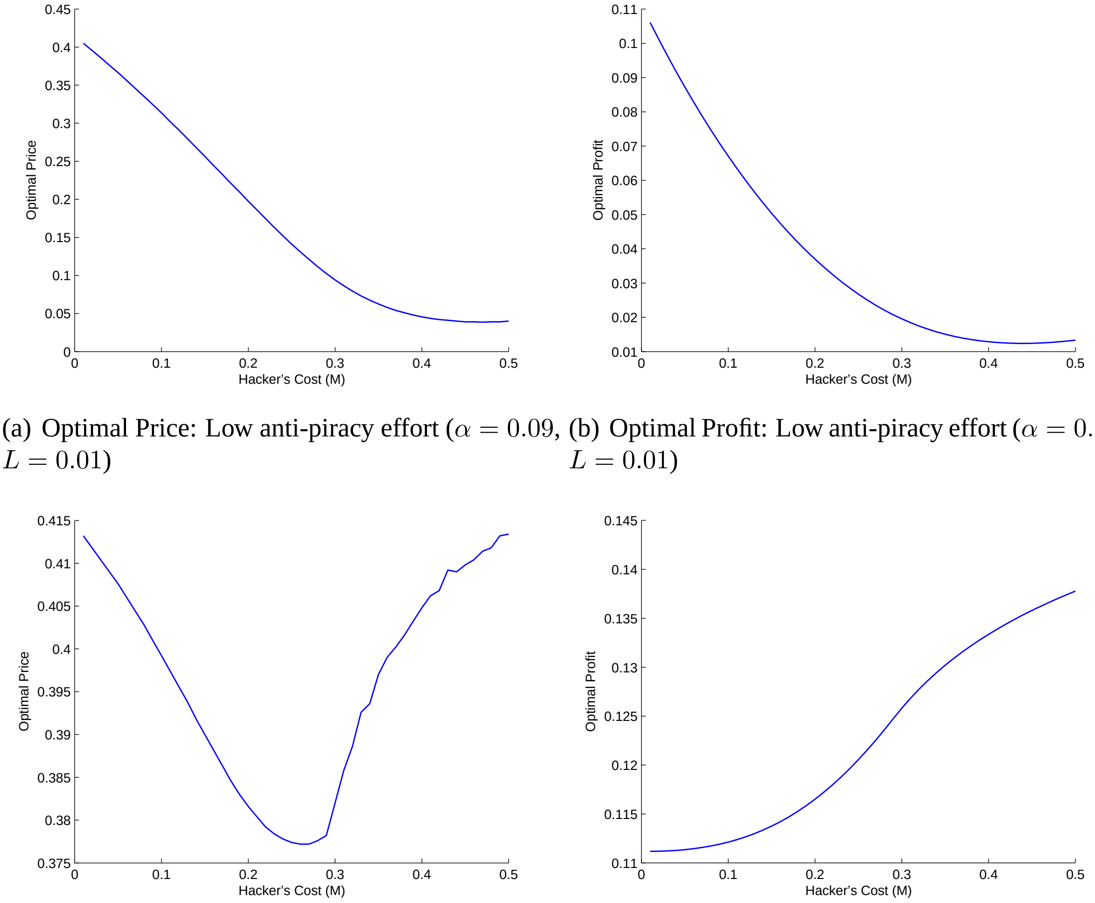

(a) Optimal Price: Low anti-piracy effort ($\alpha = 0.09$, $L = 0.01$)

(b) Optimal Profit: Low anti-piracy effort ($\alpha = 0.09$, $L = 0.01$)

(c) Optimal Price: High anti-piracy effort ($\alpha = 0.93$, $L = 0.01$)

(d) Optimal Profit: High anti-piracy effort ($\alpha = 0.93$, $L = 0.01$)

Figure 1: Effect of Hacker's cost on Optimal Price and Profit with Low Cost of Patch Quality.

the optimal vendor strategy is moderated by the cost of patch quality, L and the government's anti-piracy effort, α. We find that if L is low, then the vendor can exploit the hacker for its strategic benefit. In this case, the vendor exploits the countervailing incentive of hacking to create higher differentiation. The vendor ensures the dominance of countervailing incentive of hacking over the adverse effect of hacking by providing a high quality patch to legal users. In contrast, if L is high, the vendor does not have incentive to use the hacker strategically. This surprising result ensues from the dominance of adverse effect of hacking over the countervailing benefit of hacking when L is high.

Independent of α, when the cost of patch quality, L**, is *low***, the optimal price decreases initially with the increase in M and then increases (see Figure 1). Note that we have examined a number of cases to validate/invalidate our observations. For brevity, we have included only representative figures to demonstrate our observations. The hacker exerts a good amount of effort when M is low, and the vendor enjoys countervailing effect on $p^*_{\mathrm{Z=0}}$. Note that high hacker effort implies that users are at high risk of incurring a loss. In such a case, by providing a good patch, the vendor can increase the expected consumer surplus. This, in turn, allows the vendor to charge high price. However, as hacker's effort decreases (i.e., M increases), the attractiveness of the patch diminishes.

Even for low values of L, the vendor's profit variation due to M is influenced by α. When α is *low*, users tend to opt for the pirated version. Consequently, $p^*_{\mathrm{Z=0}}$ declines. However, at lower prices, users have the incentive to obtain the legal version; therefore, once the price declines enough, consumers opt for the legal version. Once the hacker relinquishes his activity (at high M), the vendor can focus on extracting consumer surplus without considering the impact of hacker's action. Consequently, the vendor increases the price. Figure 1(a) shows the variation in the optimal price for a low value of cost of patch quality and low anti-piracy effort level. In this case, vendor's profit also exhibits a similar pattern as price. The profit decreases due to declining price and decreasing demand. Profit bounces back for high M because of the effect of mitigated hacking activity on price.

On the other hand, if α is *high*, the vendor takes advantage of intense anti-piracy effort. In this circumstance, a declining price attracts more users to opt for the legal version. This allows the vendor to increase profit despite a declining price. With sufficiently high M, the hacker's effort becomes minimal (or non-existent); subsequently, the vendor focuses on keeping the hacker out of the market and extracting consumer surplus. This can be seen in Figure 1(c).

In contrast, if the cost of patch quality, L**, is *high***, the cost saving from decrease in the patch quality is substantial. In this case, the vendor cannot afford to provide a patch to legal users that is sufficient to ensure dominance of countervailing incentive of hacking over adverse effect of hacking in the presence of high anti-piracy effort. Not surprisingly, the vendor does not strategically benefit from hacking activity and prefers to dissuade the hacker completely (similar to the scenario where patch is available to everyone). As a result, as M increases, both the optimal price and profit increase. The following remark summarizes the aforementioned observation:

Remark 4.1. *When the patch is only provided to legal users, hacker's activity provides strategic benefit to the vendor if the cost of patch quality is low.*

4.3 Comparative Statics

In this section, we focus on comparing the outcome of the two strategies - (i) releasing the patch to everyone, i.e., $z = 1$, or (ii) providing the patch to only legal users, i.e., $z = 0$. Let us consider the impact of anti-piracy effort. As observed, anti-piracy effort has similar directional effect on the outcomes of both strategies. However, the magnitude of the effect is different for $z = 0$ and $z = 1$.

Remark 4.2. **(i)** *When the anti-piracy effort is low, the vendor provides better quality patch and charges higher price if the patch is only distributed to legal users.*

(ii) *For low α, the vendor profit is higher for $z = 0$ compare to $z = 1$.*

Remark 4.2 can be explained as follow. Recall that both α and β serve to vertically differentiate between the legal and the pirated versions, although β's effect exists only when $z = 0$. For low α, the differentiation from β is exploited, and the vendor's profits are higher. When the government exerts mild anti-piracy effort, keeping the hacking effect constant on all users, users generate relatively decent amount of surplus from the pirated version. As a result, if the patch is provided to everyone, the vendor does not benefit greatly when the anti-piracy effort is low. However, if the patch is only given to legal users, hacking activity also serves to differentiate. As mentioned earlier, hacking has adverse effect as well as countervailing effect. If the pirates have a large surplus, it is optimal for the vendor to strategically use the hacker to marginalize the pirates while protecting the legal users through a good quality patch. In this case, the marginal benefit of differentiating through hacking activity dominates the adverse effect of hacking. As such, the countervailing effect of hacking is significant. Thus, vendor's action to thwart piracy through restricting access to the patch to only legal users complements government's anti-piracy action. In other words, the vendor generates higher profit by setting $z = 0$ when the anti-piracy effort α is *low*.

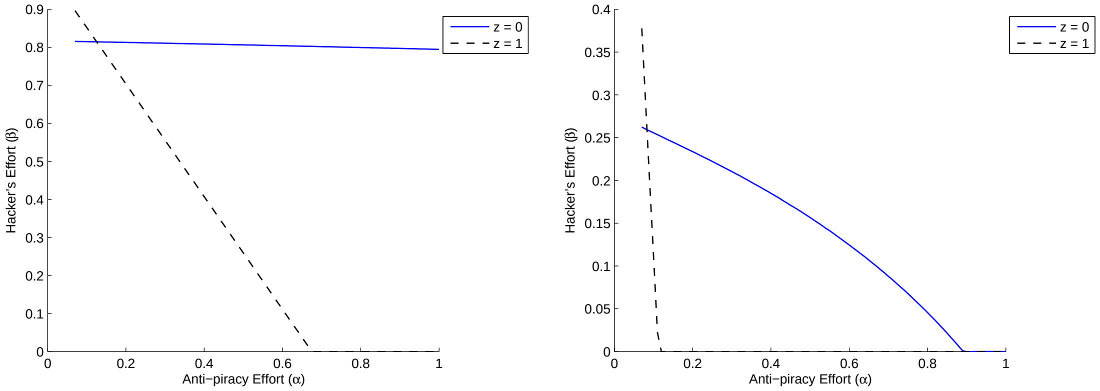

(a) Hacker's Effort ($M = 0.05$, $L = 0.01$) (b) Hacker's Effort ($M = 0.30$, $L = 0.01$)

Figure 2: Optimal Hacker's Effort - Effect of α.

It is obvious that as α increases, the vendor has lesser incentive to use β to differentiate. So, it tries to decrease the hacking activity as β has an adverse effect independent of the value of z. Since, when $z = 1$, β only destroys the consumer surplus that the vendor can extract, the vendor dissuades the hacking activity more aggressively when $z = 1$ than when $z = 0$. This can be seen in Figure 2. More formally, if $\hat{\alpha}$ denotes the anti-piracy effort level when the hacker quits hacking, then $\hat{\alpha}_{\text{z = 1}} \leq \hat{\alpha}_{\text{z = 0}}$. This is because as anti-piracy effort intensifies, when $z = 1$, all users have access to a better quality patch and the hacker gets marginalized. In contrast, if $z = 0$, pirates do not have access to the patch, which incentivizes the hacker to exert effort. Once the hacker quits, the vendor concentrates on keeping the hacker out of the market and extracting consumer surplus. As a result, for intense anti-piracy effort (α is high), the vendor takes full advantage of α as the product differentiator when $z = 1$. In contrast, for $z = 0$, in addition to the α effect there exists hacker effect. Since high α leaves negligible consumer surplus for pirated copy, the marginal benefit of hacking activity as product differentiator is insignificant. However, the vendor needs to deal with the adverse effect of hacking. Thus, the adverse effect dominates the countervailing effect of hacking for high α. Consequently, setting $z = 1$ leads to higher price and profit than setting $z = 0$ when anti-piracy effort is *high*.

Remark 4.3. (i) *There exists an* $\vec{\alpha}$ *such that* $\forall\ \alpha \in [\vec{\alpha}, 1]$, *the optimal price is higher when the patch is released to every one compare to the optimal price when patch is made available to only legal users.*

(ii) *There exists an* $\tilde{\alpha}$ *such that* $\forall\ \alpha \in [\tilde{\alpha}, 1]$, *the optimal profit is higher when the patch is released to every one compare to the optimal profit when patch is made available to only legal users.*

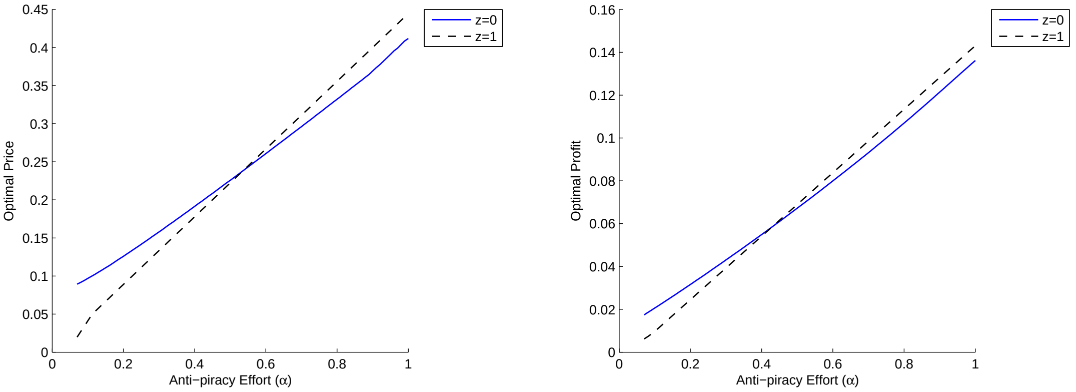

(a) Optimal Price ($M = 0.30$, $L = 0.01$)

(b) Optimal Profit ($M = 0.30$, $L = 0.01$)

Figure 3: Comparative Statics - Effect of α.

Remark 4.3 implies that, for sufficiently high α, the vendor may not benefit from vertical differentiation achieved through restricted patch distribution (see Figure 3). This is an interesting result in that quality differentiation does not create enough strategic advantage for a monopolist.

This stems from the fact that the interaction between the consumer and the vendor is influenced by actions of two other entities, namely the government and the hacker. As the hacker's cost, M, increases, the critical value $\hat{\alpha}$ decreases. This implies that the vendor will be encouraged to release patches to all users at even lower level of α when the cost of hacking increases.

5 Social Welfare

In defining social welfare, we include the net benefits of all users and the vendor. In the same spirit of Trumbull (1990), we exclude the hacker's benefit and cost from the welfare measure. Note that the payments made by the users to the vendor are transfers, and subsequently cancel out in social welfare calculation. To derive basic insights about the social welfare implications of the vendor's strategy, we define social welfare without accounting for the cost to exert effort. We will later discuss the implications of imposing the cost of effort. Thus, social welfare simplifies to,

$$\begin{aligned} \text{SW} &= (1-\beta(1-x\,z))\,(1-\alpha)\int_0^{\bar{\theta}} \theta^2\,\mathrm{d}\theta + (1-\beta(1-x))\,\int_{\bar{\theta}}^1 \theta^2\,\mathrm{d}\theta + L\log(1-x) \\ &= \frac{1}{3}\big(1\,-(\alpha+(x(z(\alpha-1)+1)-\alpha)\beta)\bar{\theta}^3+(x-1)\beta\big)+L\log(1-x). \end{aligned} \tag{9}$$

The objective of a social planner is to maximize the social welfare. In doing so, the government needs to choose an appropriate level of α or facilitate policies based on the level of α it can exert. By substituting the appropriate optimal values of p, x, and β in expression in (9), we can obtain $\text{SW}_{\text{z=0}}$ and $\text{SW}_{\text{z=1}}$, which are then compared. We observe that for low values of α, it is social welfare improving to strategically exploit hacker's action (setting $z = 0$) (see Figure 4). It is intuitive to see that, if $\alpha = 0$, there exists no market in the absence of hacking activity and restriction on patch distribution. In other words, if anti-piracy effort is extremely costly, then complementing the government's action with hacking activity through restricted patch availability is welfare improving. Nevertheless, if the anti-piracy effort is costless, then the government can

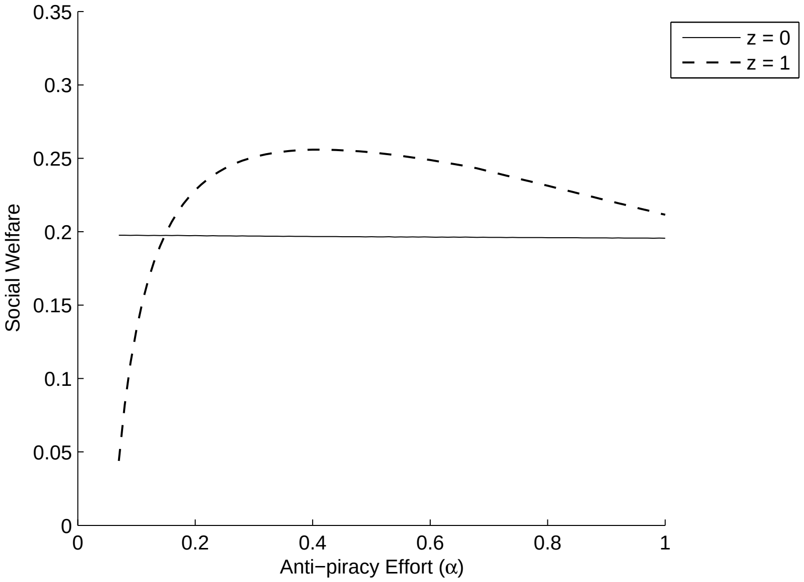

Figure 4: Social Welfare ($M = 0.05$, $L = 0.01$) .

achieve optimal welfare by choosing a moderate level of α and requiring the vendor to release the patch to everyone. Since intense anti-piracy effort reduces the surplus of pirated copy users drastically, it is not socially optimal to exert intense anti-piracy effort.

Remark 5.1. **(i)** *If it is costly to exert anti-piracy effort, then strategically exploiting hacking activity through restricting access to the patch to only legal users is social welfare improving.*

(ii) *Maximum social welfare can be obtained by exerting a moderate level of anti-piracy effort and requiring the vendor to release the patch to everyone.*

Remark 5.1 has important implications for policy makers. Strikingly, it suggests that hacking activity can be a complement of anti-piracy effort in improving social welfare. Realistically, the government may not have enough resources to exert sufficient anti-piracy effort. In such a case, the government should support restricting access to the patch to only legal users. However, if the government can afford to actively pursue anti-piracy drive, then policy makers should require software vendors to release the patch to all users and provide sufficient product differentiation through anti-piracy effort.

6 Conclusion

Lately, software vendors are taking active role in thwarting piracy by restricting access to security patches. One prominent example is Microsoft employing Windows Genuine Program to authenticate the software before downloading patches. The implications of the vendor's decision to restrict the patch only to the legal users are analyzed in this paper. In doing so, we consider roles of both the government (in the form of anti-piracy effort) and the hacker.

Our analysis compares the scenario where the vendor restricts the patch to that where the vendor offers the patch universally. We execute the comparison numerically. Fortunately, since our exogenous variables are all bounded, we are able to analyze practically the entire feasible solution space. Although we have illustrated the insights with representative values of exogenous variables, the pattern that we have observed is consistent.

Our analysis identifies two different effects of the hacker's activity. The adverse effect, which occurs independent of the vendor's decision to restrict the patch, decreases the welfare of the legal users that the monopolist can extract. In contrast, the countervailing effect occurs only when the patches are restricted. The significant negative impact of the hacker's activity on the pirates compare to that on the legal users helps the vendor vertically differentiate the legal copy from the illegal one. We observe that if the cost of the patch quality is high, then in the presence of intense anti-piracy action the vendor does not benefit from vertically differentiating between the legal and the pirated versions. This is because the countervailing incentive of hacking is negligible in such an instance; more specifically, the adverse effect of hacking dominates.

August and Tunca (2006a) find that if the population is less likely to pirate, intense anti-piracy enforcement should be complemented with the decision to restrict the patch. However, we notice that there exists a level of anti-piracy effort by the government above which vertically differentiating the legitimate copy from the pirated copy is not optimal. Rather, the vendor is better-off distributing the patch universally. In this case, complementing the government's anti-piracy effort with the vendor's decision to restrict reduces the legal user's relative willingness to pay compared to the case where only the government's anti-piracy action is utilized.

Interestingly, if the government is unable to exert anti-piracy effort (may be due to lack of resources or high cost), we find that it is social welfare improving to restrict the patch only to the legal users. This implies that for low level of anti-piracy effort, complementing it with countervailing incentive of the hacking activity is welfare improving. This result is in contrast with August and Tunca (2006a), where they find the opposite in case of low anti-piracy enforcement.

Despite the fact that the incentive of the vendor to universally provide the patch increases with the intensity in anti-piracy effort, the vendor may not undertake the socially optimal decision with respect to restricting the patch. Notwithstanding, if the government is free to choose an appropriate level of anti-piracy effort, then it is optimal to choose a moderate level of anti-piracy effort and require the vendor to release the patch to all users.

Although our results provide interesting insights regarding the implications of restricted patch distribution, our analysis is not without limitations. The key limitation of our study is the use of numerical analysis. For the analysis, we had to increment various parameters in steps of an arbitrary value (e.g., 0.01), which limited us to finite number of cases. While it does not appear that there was any discontinuity in our endogenous values over the feasible solution space, it cannot be completely ruled out. Clearly, this is a limitation of any similar numerical analysis. We also assume that if the patch is available to the user, the user will patch the system. However, it is not so in reality. Moreover, implementing the restricted patch distribution policy need not be costless. Relaxing these assumptions will lead to better understanding of the information security landscape. Several other interesting extensions are also possible. It would be insightful to generalize the model to competitive setting. One may consider that government can exert effort to make hacking more expensive. Another avenue of interesting future research would be to empirically validate our model.

References

Anderson, Ross. 2001. Why information security is hard-an economic perspective. *Proceedings of the 17th Annual Computer Security Applications Conference*. IEEE Computer Society.

Arora, Ashish, Jonathan P. Caulkins, Rahul Telang. 2006. Sell first, fix later: Impact of patching on software quality. *Management Science* **52**(3) 465–471.

Arora, Ashish, Ramayya Krishnan, Anand Nandkumar, Rahul Telang, Yubao Yang. 2004. Impact of vulnerability disclosure and patch availability - an empirical analysis. *Workshop on Economics and Information Security*. Minneapolis, MN, USA.

August, Terrence, Tunay I. Tunca. 2006a. Let the pirates patch? an economic analysis of network software security patch restrictions. *Workshop on Information Systems and Economics*. Evanston, Illinois, USA.

August, Terrence, Tunay I. Tunca. 2006b. Network software security and user incentives. *Management Science* **52**(11) 1703–1720.

Cavusoglu, Huseyin, Birendra Mishra, Srinivasan Raghunathan. 2005. The value of intrusion detection systems in information technology security architecture. *Information Systems Research* **16**(1) 28–46.

Chen, Yehning, Ivan Png. 2003. Information goods pricing and copyright enforcement: Welfare analysis. *Information Systems Research* **14**(1) 107–123.

Conner, Kathleen Reavis, Richard P. Rumelt. 1991. Software piracy: An analysis of protection strategies. *Management Science* **37**(2) 125–139.

Fudenberg, Drew, Jean Tirole. 1991. *Game Theory*. The MIT Press.

Gal-Or, Esther, Anindya Ghose. 2005. The economic incentives for sharing security information. *Information Systems Research* **16**(2) 186–208.

Gopal, Ram D., G. Lawrence Sanders. 1997. Preventive and deterrent controls for software piracy. *Journal of Management Information Systems* **13**(4) 29.

Gordon, L. A., M. P. Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security* **5**(4) 438–457.

Gordon, L. A., M. P. Loeb, W. Lucyshyn. 2003. Sharing information on computer systems security: An economic analysis. *Journal of Accounting and Public Policy* **22**(6) 461–485.

Kannan, Karthik, Rahul Telang. 2005. Market for software vulnerabilities? think again. *Management Science* **51**(5) 726–740.

Naraine, Ryan. 2005. Is mandatory windows validation a security risk? http://www.eweek.com/article2/0,1759,1755316,00.asp.

Png, I.P.L., Candy Q. Tang, Qiu-Hong Wang. 2006. Information security: User precautions and hacker targeting. National University of Singapore.

Shy, Oz. 2001. *The Economics of Network Industries*. Cambridge University Press.

Takeyama, Lisa N. 1994. The welfare implications of unauthorized reproduction of intellectual property in the presence of demand network externalities. *Journal of Industrial Economics* **42**(2) 155.

Trumbull, William. 1990. Who has standing in cost-benefit analysis? *Journal of Policy Analysis and Management* **9**(2) 201–218.

Varian, Hal R. 2000. Buying, sharing and renting information goods. *Journal of Industrial Economics* **48**(4) 473–488.