Information Systems Research
Vol. 19, No. 1, March 2008, pp. 48–70
ISSN 1047-7047 | EISSN 1526-5536 | 08 | 1901 | 0048

informs

DOI 10.1287/isre.1070.0142

Let the Pirates Patch? An Economic Analysis of Software Security Patch Restrictions

Terrence August
Rady School of Management, University of California at San Diego, La Jolla, California 92093, taugust@ucsd.edu

Tunay I. Tunca
Graduate School of Business, Stanford University, Stanford, California 94305, tunca_tunay@gsb.stanford.edu

We study the question of whether a software vendor should allow users of unlicensed (pirated) copies of a software product to apply security patches. We present a joint model of network software security and software piracy and contrast two policies that a software vendor can enforce: (i) restriction of security patches only to legitimate users or (ii) provision of access to security patches to all users whether their copies are licensed or not. We find that when the software security risk is high and the piracy enforcement level is low, or when tendency for piracy in the consumer population is high, it is optimal for the vendor to restrict unlicensed users from applying security patches. When piracy tendency in the consumer population is low, applying software security patch restrictions is optimal for the vendor only when the piracy enforcement level is high. If patching costs are sufficiently low, however, an unrestricted patch release policy maximizes vendor profits. We also show that the vendor can use security patch restrictions as a substitute to investment in software security, and this effect can significantly reduce welfare. Furthermore, in certain cases, increased piracy enforcement levels can actually hurt vendor profits. We also show that governments can increase social surplus and intellectual property protection simultaneously by increasing piracy enforcement and utilizing the strategic interaction of piracy patch restrictions and network security. Finally, we demonstrate that, although unrestricted patching can maximize welfare when the piracy enforcement level is low, contrary to what one might expect, when the piracy enforcement level is high, restricting security patches only to licensed users can be socially optimal.

Key words: IT security; software piracy; IT policy and management; network economics; economics of IS
History: Sanjeev Dewan, Senior Editor; Rahul Telang, Associate Editor. This paper was received on May 9, 2006, and was with the authors $3\frac{1}{2}$ months for 3 revisions.

1. Introduction

Piracy has long been an important concern for the software industry. The relative ease of replicating and distributing software, as with any digital good, combined with the high value that many software products command makes software a prime target for piracy and unlicensed use. Today, an estimated every third copy of Microsoft's widely used Windows operating system is unlicensed (Fried 2005), and the ratio of pirated software reaches up to 90% of total usage in certain countries (BSA-IDC 2005). The estimated total cost of software piracy and counterfeiting in the United States alone is about $7 billion per year, and the annual global cost of piracy exceeds $30 billion (Rooney 2005). Furthermore, in the contemporary global technology environment characterized by broad Internet connectivity and frequent threats on network security, the impact of software piracy on vendors is not simply limited to lost revenues from unrealized sales. Rather, unprecedented new issues and challenges such as security interdependence among interconnected systems and the related incentive problems in a network environment arise and pose new complications brought about by software piracy.

To see the highly challenging nature of the problem, consider the dilemma recently experienced by Microsoft. On January 26, 2005, Microsoft announced that, as part of its newly implemented "Genuine Advantage" program introduced to combat piracy, it would require users to validate their individual copies of the Windows XP operating system before

permitting them to download updates (such as the essential Service Pack 2 (SP2) update, which was released in August 2004). The most critical impact of this decision was that users of pirated copies of Windows XP would not be able to patch their systems with security updates (Microsoft 2005a). This dimension of Microsoft's decision stirred a great deal of discussion among IT security experts and the broader community on whether Microsoft was jeopardizing Internet security as well as hurting its own profitability by taking such a stance. From the company's own point of view, the decision is not an easy one, and there is a complicated trade-off that needs to be addressed: On one hand, opponents of the decision point out that such a restriction would significantly compromise the security of the network by creating a large population of "unpatched hosts" on the Internet, which are susceptible to "infection" and can spread malicious code such as worms and viruses (see, e.g., Moore et al. 2002, Weaver et al. 2003, and Schneier 2004 among others). Under the high security risks faced today, such a policy reduces the security of the entire Internet, including the systems of legitimate users. As a result, in addition to facing public pressure for selfish behavior, the value of Microsoft's product is reduced, which also ultimately hurts the company (see, e.g., Hou 2004, August and Tunca 2006 and the references therein). On the other hand, proponents of the decision defend Microsoft's rights for intellectual property while stressing that this approach could help curb software piracy (see, e.g., Rooney 2005). On the heels of this hot debate, just six months after the first announcement in January, at the worldwide launch of Windows Genuine Advantage, Microsoft announced that it had changed its previously declared policy, deciding to continue to allow pirates to download security patches while restricting access to standard updates only to legitimate users (Microsoft 2005b).

Microsoft's change of policy and apparent indecisiveness is not an isolated occurrence but, rather, is a part of an ongoing dilemma for the company. Microsoft had announced security patch restrictions and reversed such decisions before (Microsoft 2002, Worthington 2004, Salloway 2004) and at times applied such restrictions selectively for certain cases and countries among other measures to combat piracy (Evers 2005, Bass 2005). At the heart of the issue lies the difficult decision that not only Microsoft but also any software vendor faces. Prohibiting users of pirated copies of the software from patching the product decreases the security of the network for everyone, which may have costly consequences with losses reaching up to billions of dollars every year. This decrease in software security reduces the value of the product for the buyers and decreases the vendor's sales and profit. However, in addition to punishing the users who infringe upon its intellectual copyright, restricting pirates from applying security patches can be a strategic tool for the vendor: Not allowing software pirates to download security updates and patch their systems puts them in a compromised position as they face the risks of being exposed to malicious attacks. Thus, these restrictions can increase the attractiveness of purchasing the software relative to committing piracy. As a result, a significant percentage of (would-be) pirates may elect to purchase the software, which can substantially increase vendor profits. Given this trade-off and depending on the product and market conditions, the vendor is facing a complicated policy decision on whether to allow software pirates to install security patches or to restrict such patches only to legitimate users. Furthermore, this critical policy decision has important consequences on the value and social welfare generated by the product, and, therefore, it is an important issue for governments and social policy makers as well.

These observations motivate a formal study of the economics of a vendor's security patch restriction policy decision. In this paper we aim to provide insights into the economics of a vendor's patch release policy under software piracy in connection with current empirical observations and the ongoing debate. Building on the model given in August and Tunca (2006), we explore the implications of the two alternative policies: (i) restricting the security patches only to legitimate users or (ii) providing access to security patches to all users without checking the legitimacy of their copies of the software. Our analysis has two main purposes. First, we identify the conditions under which each policy will be optimal for a software vendor. Second, we explore the implications of patch restrictions on security of a software product, piracy enforcement, and social welfare.

We present a joint model of piracy and negative network security externalities and show that when both factors are considered, depending on the vendor's pricing and patch restriction policies, a variety of consumer market structures can be induced in equilibrium. The equilibrium market structure, in turn, affects vendor profits and social welfare. Exploring the optimal patch restriction policy for the vendor, we find that, when software is highly risky or the population's tendency for piracy is high, it is optimal for the vendor to impose security patch restrictions on unlicensed users, whereas if the patching costs are sufficiently low, the profit-maximizing policy for the vendor is to allow all users, licensed or unlicensed, to apply security patches. When the population's tendency for piracy is low, the optimality of patch restrictions is contingent upon the piracy enforcement level. If the piracy enforcement level is high, a software vendor should restrict security patches only to licensed users, whereas in an environment with a low level of piracy enforcement, he should employ an unrestricted patch release policy.

Next, in the presence of software security patch restrictions, a vendor may prefer a less secure product and hence can have reduced incentives to invest in improving software security. As a result, social welfare can suffer significantly. In addition, contrary to what one may expect, we show that an increased piracy enforcement level does not always increase vendor profits. Furthermore, we show that, for certain piracy enforcement levels, governments can, in fact, increase social surplus generated by the software product by increasing piracy enforcement, thereby inducing the vendor to lower his price strategically to target pirates' incentives to convert into purchasers under high security risk. Finally, we show that, contrary to some arguments made in the software community, policies that restrict unlicensed users from patching can *increase* social welfare. In fact, we demonstrate that having the government impose laws to ensure such restrictions can sometimes be necessary to maximize the surplus generated by the software.

The rest of the paper is organized as follows. Section 2 presents a review of the relevant literature. Section 3 presents the model and studies consumer market equilibria under the two proposed patch permission policies. Section 4 explores the conditions under which each policy will be optimal. Section 5 presents our results about the effects of patch restriction policies on software security, the role of piracy enforcement, and welfare. Section 6 offers our concluding remarks. All proofs are given in the appendix at the end of this paper and in the online supplement.[1]

2. Literature Review

There is a growing literature on economics of software security. Although the research subject is relatively new, several streams of research exist in the IT literature. Anderson and Moore (2006) provide a broad overview of the existing research on economics of information security. An important stream focuses on software vulnerability disclosure. Cavusoglu et al. (2004a) examine policies for vendor vulnerability disclosure, specifically looking at full vendor disclosure, immediate public disclosure, and hybrid policies. They find that vulnerability characteristics, cost structure, and a vendor's patch development incentives determine which policy is optimal. Arora et al. (2005) investigate the optimal timing for disclosure and find that a software vendor's patch release time lags behind the social optimum. Jaisingh and Li (2005) examine how to use disclosure as a scheme to coordinate timing of patch releases. Other topics examined in the IT security literature include determination of optimal frequency of patching to balance the operational and damage costs associated with security vulnerabilities (Cavusoglu et al. 2004b), vendor incentives to invest in software quality (Arora et al. 2006), and the value of intrusion detection systems in IT security architecture (Cavusoglu et al. 2005). A more detailed survey on the information security literature and other related papers can also be found in August and Tunca (2006), with which our current paper shares its base model.

August and Tunca (2006) compare policies that target user incentives from the point of view of a profit-maximizing vendor and a social welfare-maximizing planner. Each user has her own independent system on which she makes usage and patching decisions considering the value that would be obtained

[1] An online supplement to this paper is available on the *Information Systems Research* website (http://isr.pubs.informs.org/ecompanion.html).

by employing the product, the risk associated with usage, the costs of patching, and the incentives offered by the vendor and the social planner. Unpatched systems on the network cause negative security externalities on other users. August and Tunca (2006) show that, from the points of view of both a social planner and a profit-maximizing vendor, mandating patching on the users is unlikely to be effective. On the other hand, for proprietary software, creating positive incentives for patching by users can be effective, whereas imposing usage taxes will hurt both vendor profits and social welfare. Our current paper pursues an understanding of the economic links between piracy and network security risk by combining the two factors and expanding in several critical directions. First, we introduce heterogeneity in tendencies toward piracy into the consumer population and incorporate endogenous piracy usage decisions. Second, we explore an additional control lever for the vendor beyond price, namely a policy decision of restricting security updates only to legitimate copies of his product. Third, we introduce additional market characteristics such as piracy enforcement and piracy potential into the model. With these additional model elements, the current paper combines the economics of software security and piracy and studies the resulting economic interaction. The implications of the combined model go well beyond the separate implications of piracy and software security. We identify the underlying incentives for pirated usage and the incentives that drive potential pirates toward purchasing legitimate products. We also characterize the critical impact of the market characteristics on the vendor's policy decision and explore the effects of patch restriction policies on incentives for investment in software security, on the role of piracy enforcement, and on social welfare. Our results show that many significant new economic insights and important policy implications arise from the interaction between network software security and software piracy.

Digital piracy has long been discussed by researchers in the literature. However, there have not been any studies that examine the implications of software piracy on Internet security yet. Our work bridges these two main branches of contemporary IT economics literature at an important current junction. Comprehensive surveys of economics literature on piracy of digital goods can be found in Peitz and Waelbroeck (2003b) and Varian (2005). Rao (2003) presents an overview of copyright issues for e-information. Holsing and Yen (1999) provide an overview of ethical, technical, managerial, and economic issues related to software piracy.

The traditional view on piracy is that it reduces a legal publisher's profits (Novos and Waldman 1984, Johnson 1985). Consequently, one stream of literature examines policies to combat piracy (see, e.g., Gopal and Sanders 1997, Harbaugh and Khemka 2001). Alvisi et al. (2002) suggest that piracy can cause a legal publisher to employ quality differentiation in its products in order to divert consumers from pirating the product to purchasing it. Sundararajan (2004) analyzes the joint choice of price discrimination and technological protection through Digital Rights Management, finding that the two can act as substitutes. Chen and Png (2003) explore the relative effectiveness of piracy control measures on increasing social welfare, finding that increasing detection is more detrimental to welfare than price cuts. They also show that a subsidy is optimal from the point of view of welfare maximization whereas a tax on the copying medium is welfare superior to fines for piracy. Our paper establishes that, in the presence of security threats, a policy that restricts security updates can also be used as a measure to curb piracy and, in certain cases, increase welfare. However, our results show that such a measure is not always preferable from the vendor's point of view.

Despite the traditional view on the negative effects of piracy, there is a sizable stream of literature in economics that discusses a variety of potential beneficial side effects of copyright violations for a legal publisher from many different aspects. A legal publisher can recoup benefits through indirect appropriation of the value of the pirated copies from the first purchaser (Liebowitz 1985, Besen and Kirby 1989), through the sale of ancillary products that accommodate the functioning of these products such as CDs and radio (Curien et al. 2004), or through indirect effects of file sharing on consumers such as learning (Bakos et al. 1999, Varian 2000, Peitz and Waelbroeck 2003a). A number of papers argue that, under network effects, a certain degree of piracy can benefit a legal publisher (see, e.g., Connor and Rumelt 1991; Takeyama 1994, 1997; Slive and Bernhardt 1998; Shy

Figure 1 Model Timeline for Vendor Pricing and Policy Decisions as Well as Consumer Decision Making

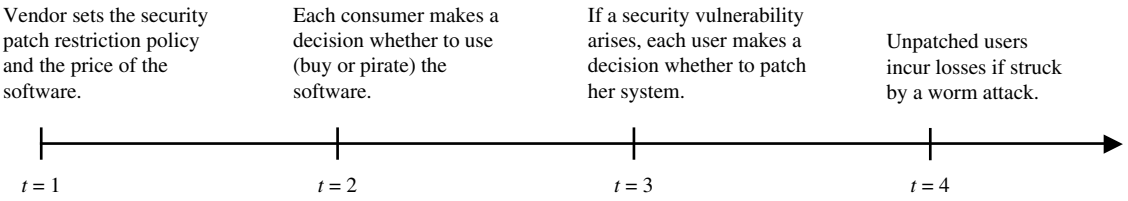

and Thisse 1999). The main argument in these studies is that allowing certain users to pirate increases the appeal of the product because of the existence of positive network effects. This effect, in turn, increases the publisher's profits from legal users. In our model we examine negative network externalities associated with piracy, in particular negative security externalities. Unlike the papers that demonstrate the benefits of piracy under positive network effects (or studies that highlight potential side effects of piracy), piracy does not directly benefit a legal publisher under negative security externalities. Yet we show that, under certain conditions, a legal publisher can find it preferable not to restrict benefits to pirates in order to limit the harms that pirates can impose on legal users.

3. The Model and Consumer Market Equilibrium

3.1. Model Description

We build on the model of August and Tunca (2006). A software vendor produces and supports a network software product. There is a continuum of consumers whose valuations for the product lie uniformly on $\mathcal{V} = [0, 1]$. The software is not perfectly secure: If it has a vulnerability, the vendor releases a patch and the consumers who purchased the product may undergo costly patching to prevent security attacks and breaches.

We add the possibility of software piracy to the structure of August and Tunca (2006). Generally, consumer tendencies toward piracy are heterogeneous. Whereas certain consumers have a low tendency for piracy or tend to behave "ethically," there are certain consumers who have a high tendency for piracy or can behave "unethically" and choose to pirate the product if the benefit from pirating exceeds its risk. To capture this variation, we model two types of consumers in the market, each consumer being one of two types in $\Theta = \{L, H\}$, where Type L denotes a consumer of "Low piracy tendency" (or no tendency in our case for simplicity) and Type H denotes a consumer of "High piracy tendency."[2] For any given consumer, the probability that she is of Type H is given by $\nu \in [0, 1]$.

There are four time periods, as illustrated in Figure 1. In the first period, the vendor sets a policy in regard to security patch restrictions for unlicensed users $\rho \in \{l, nl\}$ and sets the price p of the software. Under policy "l" the vendor "lets" unlicensed users (pirates) patch by making software patches generally available to all users in case a vulnerability arises. Under policy "nl" the vendor does "not let" pirates patch. In this case, the vendor restricts the availability of security patches only to legitimate consumers.

In the second period, given both the price and the security patch restriction policy and depending on her type, each consumer makes a decision whether to pirate the software, purchase the software, or simply not become a user. We denote these actions with S, B, and NU, respectively. If she chooses to pirate the software, she will be detected with probability $\pi_d > 0$, in which case she incurs a loss of $c_d > 0$. In the third period, it is revealed whether the software has a security vulnerability or not. If a security vulnerability exists, a software patch is made available by the vendor to all legitimate users, but, depending on the vendor's policy, it may or may not be made available to unlicensed users. Subsequently, each consumer who is permitted to patch under the vendor's policy makes a decision whether to patch her installation, trading off the risks associated with not patching

[2] Although these labels are used in the literature, one should keep in mind that differences in tendency toward piracy equivalently arise due to many reasons such as technological capability and usage context (see, e.g., Chen and Png 2003, Peitz and Waelbroeck 2003b).

versus the costs associated with patching. We denote the consumer's patching decision by P and NP, referring to "patch" or "not patch," respectively. If a consumer decides to patch her system, she incurs a cost of patching, which accounts for the money and effort that she must exert in order to verify, test, and roll out patched versions of existing systems (Bloor 2003). Taking the probability that a need for patching arises into account, we denote the expected cost of patching a system by $c_p > 0$.

Finally, if a security vulnerability arose during the second period and users made patching decisions, then a malicious attack may occur in the fourth period. If such an attack occurs, unpatched consumers may get hit and incur losses. We denote the probability that there is a security vulnerability and a security attack on the network as $\pi_a > 0$. If the mass of the unpatched population in the network is u, then the probability that the attack will successfully penetrate the network and hit an unpatched user is $\pi_a u$. As in August and Tunca (2006), if a user with valuation v goes unpatched and is hit by the attack, the loss that she suffers is αv, where $\alpha > 0$ is a constant.

For each ρ, the action space for Type L consumers is $S_\rho^L = \{B, NU\} \times \{P, NP\} - (NU, P)$. The exclusion of (NU, P) arises due to infeasibility. In a similar manner, the action spaces for Type H consumers under $\rho = l$ and $\rho = nl$ are $S_l^H = \{B, S, NU\} \times \{P, NP\} - (NU, P)$ and $S_{nl}^H = S_l^H - (S, P)$, respectively. Under $\rho = nl$, (S, P) is not feasible because the vendor has restricted pirates from receiving security patches. Given the price p, security patch restriction policy ρ, expected loss when committing piracy $\pi_d c_d$, expected cost of patching c_p, and effective security risk $\pi_a \alpha$, in a consumer market equilibrium, each consumer maximizes her expected utility taking the equilibrium strategies for all consumers as fixed. For a strategy profile σ: $\mathcal{V} \times \Theta \to \bigcup_{\theta \in \Theta} S_\rho^\theta$, when a security vulnerability is revealed, the expected security cost faced by the consumer with valuation v and type θ is then defined by

$$C(v, \theta, \sigma) \triangleq \begin{cases} \pi_a \alpha u(\sigma) v & \text{if } \sigma(v, \theta) \in \{(B, NP), (S, NP)\}; \\ c_p & \text{if } \sigma(v, \theta) \in \{(B, P), (S, P)\}; \\ 0 & \text{if } \sigma(v, \theta) \in \{(NU, NP)\}, \end{cases} \quad (1)$$

where $u(\sigma) \triangleq \int_{\mathcal{V}} 1_{\{\sigma(v, \theta) \in \{(B, NP), (S, NP)\}\}} \, dv$.[3] The expected cost of usage for the consumer with valuation v and type θ is given by

$$P(v, \theta, p) \triangleq \begin{cases} p & \text{if } \sigma(v, \theta) \in \{(B, NP), (B, P)\}; \\ \pi_d c_d & \text{if } \sigma(v, \theta) \in \{(S, NP), (S, P)\}; \\ 0 & \text{if } \sigma(v, \theta) \in \{(NU, NP)\}. \end{cases} \quad (2)$$

Note that both ρ and θ affect the strategy sets of the consumers and consequently the applicable region in Equations (1) and (2). The surplus gained by consumer (v, θ) by employing the software will then be $v - C(v, \theta, \sigma) - P(v, \theta, p)$. The consumers who do not patch cause a negative externality on all users by decreasing the safety of the network and the software. Clearly, for any $v \in \mathcal{V}$, $C(v, \theta, \sigma)$ defined by (1) is increasing in $u(\sigma)$ (i.e., the unpatched population). Furthermore, consumers who patch protect themselves from the negative externality caused by the unpatched population. The exogenous parameter space of the model is α, $c_d \in (0, \infty)$ and c_p, π_a, π_d, $\nu \in (0, 1)$.

Importantly, note that the factors ν and $\pi_d c_d$ measure very different aspects of piracy. In every country and consumer market for a digital good, the tendency for piracy varies substantially across the consumers. In addition to standard deterrents of piracy such as legal penalties, depending on idiosyncratic factors such as technological capability, socioeconomic position, ethical considerations, and nature of business, both consumer aversion to committing piracy and the probability of being detected vary. Consequently, each consumer has a predisposed tendency for piracy, which arises as a combination of market and private factors and determines her specific "cost" of pirating the product. We capture this variation in cost of piracy across consumers by considering the two types (i.e., L or H) that each consumer can belong to which determine her cost level or tendency for piracy. From this point of view, Type L consumers can be thought of as having a very high expected cost of piracy, whereas Type H consumers can be thought of as having a relatively low cost of piracy. The parameter ν captures this dimension of variability in consumer likeliness to

[3] The notation "$\triangleq$" has the meaning "as a definition" throughout the paper.

pirate the product, whereas $\pi_d c_d$ captures the "base" level of punishment and enforcement or expected cost for piracy.[4] As such, ν is the exogenous determinant of how likely it is for consumers in the market to be predisposed to piracy, and, for convenience in terminology, we will refer to it as the population's *tendency for piracy* throughout the paper. On the other hand, once a consumer has a high (economic) tendency to pirate, $\pi_d c_d$ factors in the endogenous decision she makes whether to actually pirate the product, and we refer to it as the *piracy enforcement level* throughout the paper.

3.2. Consumer Market Equilibrium

Before addressing the vendor's profit maximization problem, we must first determine the consumer market equilibrium for any given price p and the patch restriction policy ρ. Each consumer with valuation v and type θ chooses the action in the strategy set S_ρ^θ (which depends on both her type and the restriction policy) that maximizes her net payoff. As a result, her optimal action is determined by solving the maximization problem

$$\begin{aligned} \max_a \quad & v \cdot 1_{\{a \neq (NU, NP)\}} - C(v, \theta, \sigma_{-v}) - P(v, \theta, p) \\ \text{s.t.} \quad & a \in S_\rho^\theta, \end{aligned} \tag{3}$$

where σ_{-v} indicates that all other consumers' strategies are held fixed. The solution of (3) for each $v \in \mathcal{V}$ and $\theta \in \Theta$, namely the function $a^*(v, \theta)$, gives rise to an equilibrium strategy profile σ^*. This equilibrium strategy profile describes which consumers are foregoing, pirating, and purchasing the software as well as which consumers are patching and not patching their respective copies of the software. Next, we provide a full characterization of the equilibrium strategy profile.

When the price is sufficiently low (i.e., $p \leq \pi_d c_d$), none of the consumers who can pirate optimally do so in equilibrium because the price is lower than the expected loss when committing piracy. Therefore, for $p \leq \pi_d c_d$, there is no piracy in equilibrium, and the characterization of the consumer market equilibrium is identical to that given in Lemma 1 in August and Tunca (2006). A consumer (of Type L or H) will buy the product if her valuation is sufficiently high enough to absorb the potential expected losses in case of an attack in addition to the price, i.e., if $v \geq p + \pi_a \alpha u(\sigma^*) v$. On the other hand, a consumer will patch in case of a vulnerability if the potential security risk she faces exceeds the patching costs, i.e., if $c_p \leq \pi_a \alpha u(\sigma^*) v$. The resulting equilibrium is characterized by two thresholds: a consumer with valuation v will buy the product if $v \geq v_b$ and patch if $v \geq v_p$, where $p < v_b \leq v_p \leq 1$. This characterization is consistent with what is typically seen in business and personal software usage. Corporate or high-valuation consumers tend to be the ones performing responsible patch management whereas lower-valuation or personal users are more likely to go unpatched and bear the security risks.[5]

When $p \leq \pi_d c_d$, the patch restriction policy of the vendor does not play a role in shaping the equilibrium because it is not rational for Type H consumers to pirate the software product. However, when $p > \pi_d c_d$, piracy becomes potentially attractive for Type H consumers, and the consumer equilibrium behavior will be affected by the patch restriction policy ρ that the vendor employs. Thus, ρ has an important effect on the consumer equilibrium behavior, and the resulting market structures need to be characterized contingent on ρ for $p > \pi_d c_d$.

3.2.1. No Security Patch Restrictions ($\rho = l$). Consider the policy l under which the vendor allows pirates to patch their systems should security vulnerabilities arise. In this case, the consumers who choose to pirate the product face no downside from security patch restrictions, and, given $p > \pi_d c_d$, they may have incentives to pirate the product. As a result, the consumer market structure changes with a shift to piracy. Lemma A.1 in the appendix presents the equilibrium for this case, which is characterized again by buying and patching thresholds on consumer valuations. For Type L consumers, the threshold valuations v_b and v_p (although they may be different) play

[4] Equivalently, including other private costs of piracy, one can think of Type H consumers having an expected cost of $\pi_d c_d$ for committing piracy and Type L consumers having an expected cost of $\pi_d c_d + K$, where K is prohibitively high to preclude piracy activity by such consumers.

[5] See August and Tunca (2006) for further discussion.

Figure 2 Two Possible Equilibrium Consumer Market Structures Under the Policy of Letting Pirates Patch (i.e., $\rho = l$) with Cutoff Consumer Valuations as Defined in Lemma A.1.

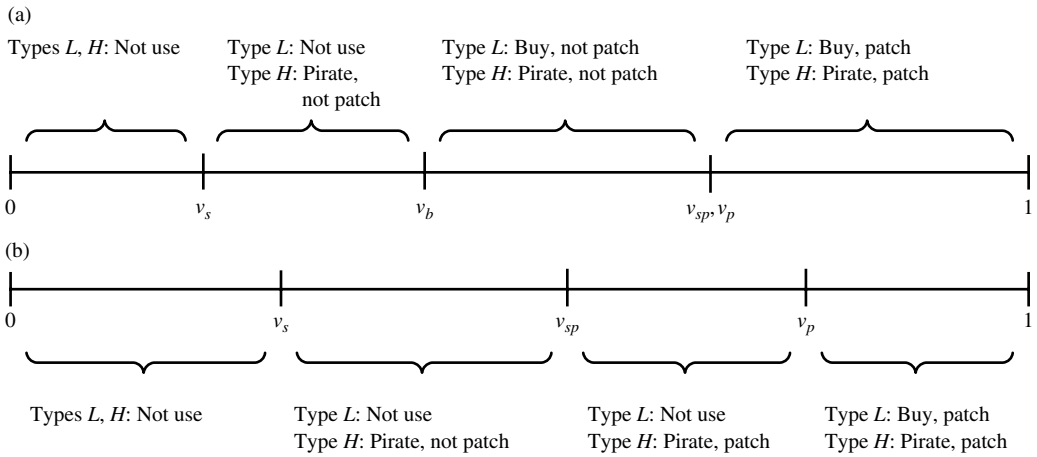

Notes. Panel (a) presents a structure in which some Type *L* consumers purchase but do not patch. Panel (b) presents a structure in which all purchasing Type *L* consumers patch.

the same role as described above. However, Type H consumers pirate the product as opposed to purchasing it because their expected cost from pirating is less than the cost of purchasing the product (i.e., $p > \pi_d c_d$). A Type H consumer with valuation v will use the product if and only if $v \geq \pi_d c_d + \pi_a \alpha u(\sigma^*)v$, and she will patch if and only if $c_p \leq \pi_a \alpha u(\sigma^*)v$. The resulting pirating and patching thresholds for Type H customers are denoted by v_s and v_{sp}. That is, Type H consumers with valuations lower than v_s are non-users, those with valuations between v_s and v_{sp} pirate but do not patch, and, finally, those with valuations greater than v_{sp} both pirate and patch the software.

Depending on the relative order of the cutoff valuations given in Lemma A.1, the consumer market structure can take on several different forms. In each market structure, which consumers are patching, buying, pirating, and not using as determined by their valuation and type can vary significantly. Figure 2 demonstrates two of the possible market structures under $\rho = l$. Panel (a) demonstrates a case where there are both patching and non-patching pirates and legitimate purchasers. Panel (b) demonstrates a case where there are both patching and non-patching pirates while all legitimate purchasers of the software opt to patch the product. Panel (a) corresponds to a case where effective security risk ($\pi_a \alpha$) is relatively low compared to that in panel (b).[6]

3.2.2. Security Patch Restrictions for Unlicensed Users ($\rho = nl$). When the vendor restricts the availability of the security patches (i.e., when $\rho = nl$), Type H consumers face an additional trade-off. Because a Type H consumer cannot apply a security patch if she chooses to pirate the product, she faces a security risk in case a vulnerability arises. Instead, she may prefer to purchase the product in order to gain the right to apply security patches. As a result, when $\rho = nl$, there may be Type H consumers who are purchasers (and patchers) in equilibrium, even though $p > \pi_d c_d$. The resulting equilibrium is given in Lemma A.2 in the appendix. As for the case for $\rho = l$, the equilibrium is again characterized by four threshold valuations for $\rho = nl$. However, there are important differences between the consumer market equilibrium characterizations under the two policies. First, under the policy nl, there may be Type H consumers who could pirate the product but choose to buy instead. High-value Type H consumers (i.e., those

[6] Depending on the parameters and the price p, under the policy l, there are two additional possible market structures. Full mathematical characterizations of all four market structures as well as those for the threshold valuations are provided in Lemmas B.1 and B.2 in the online supplement.

Figure 3 Two Possible Equilibrium Consumer Market Structures Under the Policy of Not Letting Pirates Patch (i.e., $\rho = nl$) with Cutoff Consumer Valuations as Defined in Lemma A.2.

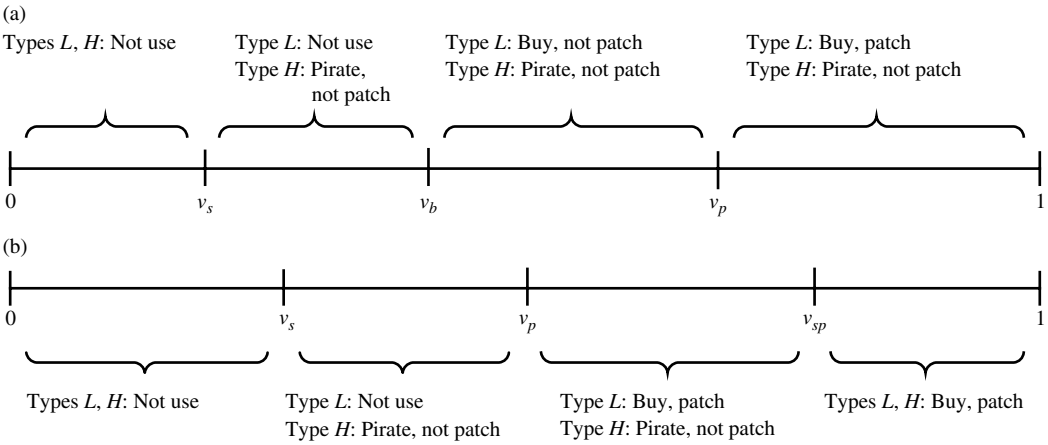

Notes. Panel (a) presents a structure in which no Type *H* consumer purchases. Panel (b) presents a structure in which some Type *H* consumers purchase.

with valuations greater than v_{sp}) now *convert* into buyers who purchase and patch the software. Specifically, the Type H consumer with valuation v will convert to a legitimate patching user if and only if her valuation satisfies $p - \pi_d c_d + c_p \leq \pi_a \alpha u(\sigma^*)v$. Second, when pirates are not allowed to patch, the cutoff patch valuation for Type H consumers (v_{sp}) is always greater than or equal to the cutoff patch valuation for Type L consumers (v_p). Figure 3 presents two possible equilibrium market structures under the policy $\rho = nl$. In the market structure presented in panel (a), no pirates are converted into buyers. There are some Type H consumers who elect to pirate the product but not patch as well as some Type L consumers who buy the product but choose not to patch. In addition, there are some Type L consumers who both buy and patch the product. Such an equilibrium market structure would arise in a case where the price is high and the effective security risk is relatively low. In the market structure shown in panel (b), some would-be pirates are converted into buyers in equilibrium because Type H consumers who have valuations higher than v_{sp} choose to buy and patch the software. Some Type H consumers with lower valuations choose to pirate the software but cannot patch because of the restrictions. However, they generate sufficient negative security externalities resulting in a market structure where all buying Type L consumers patch in case of a vulnerability. This type of market structure can arise when price is relatively low but security risk is relatively high.[7]

4. When Is It Optimal for a Vendor to Impose Security Patch Restrictions?

We start our analysis by examining the question when should a vendor impose security patch restrictions. We first present the vendor's decision making problem. Then, we explore the market and product characteristics under which the vendor should optimally restrict security patches to legitimate users only and under which he should let the pirated copies of his software be patched.

4.1. The Vendor's Profit-Maximization Problem

The software vendor faces two decisions he must make to maximize his profits in this market environment with piracy. First, he chooses a policy $\rho \in \{l, nl\}$, which determines whether pirates will be allowed to patch their software if a vulnerability arises. Second, he must set the consumer price. Without loss of generality, we assume that the marginal cost of production for each copy of the software is zero. When the

[7] There are four additional possible equilibrium market structures for $\rho = nl$. The complete characterization of all six market structures together with the characterizations for the threshold valuations given in Lemma A.2 are provided in Lemmas B.1 and B.3 in the online supplement.

vendor permits pirates to patch their software installations (i.e., $\rho = l$), his expected profit function is defined by

$$\Pi_l(p) \triangleq \begin{cases} p(1-v_b) & \text{if } p \le \pi_d c_d; \\ p(1-\nu)(1-v_b) & \text{if } p > \pi_d c_d. \end{cases} \quad (4)$$

On the other hand, when the vendor restricts patches to only the legitimate users (i.e., $\rho = nl$), his expected profit function is defined by

$$\Pi_{nl}(p) \triangleq \begin{cases} p(1-v_b) & \text{if } p \le \pi_d c_d; \\ p(\nu(1-v_{sp}) + (1-\nu)(1-v_b)) & \\ \qquad \text{if } p > \pi_d c_d. \end{cases} \quad (5)$$

Both profit functions reflect that, when $p \le \pi_d c_d$, both types of consumers who use the software are purchasers, because no user will have an incentive to pirate in this case. However, for $p > \pi_d c_d$, the choice of a security patch restriction policy becomes quite relevant. As can be seen in (4), when $\rho = l$, the main concern for the vendor is how Type L consumers value his software product. On the other hand, when $\rho = nl$, the vendor makes profits not only from Type L consumers but also from Type H consumers who may choose to purchase the product, as shown in (5). We define p_l^* and p_{nl}^* as the prices that maximize $\Pi_l(\cdot)$ and $\Pi_{nl}(\cdot)$, respectively. Given the effective patching cost, effective security risk, expected loss when committing piracy, probability of being a Type H consumer, and consumer market equilibria established in Lemmas A.1 and A.2, the vendor's problem can be written as

$$\begin{aligned} \max_{\rho,\, p} \ & \Pi_\rho(p) \\ \text{s.t. } & 0 \le p \le 1 \\ & \rho \in \{l, nl\}. \end{aligned} \quad (6)$$

Note that both p and ρ have a significant effect on the threshold valuation levels (v_s, v_{sp}, v_b, and v_p) that characterize the consumer market structures that arise. Because both $\Pi_l(\cdot)$ and $\Pi_{nl}(\cdot)$ are bounded on a compact domain and have a single discontinuity at which they are left-continuous and decreasing to the right, there exists a solution to (6). We denote the solution to the vendor's problem with (ρ^*, p^*).

4.2. When Is Restricting Pirates from Patching Optimal?

We start by exploring the market conditions where it can be optimal for the vendor to restrict unlicensed users from applying security patches. The following proposition presents the result.[8]

Proposition 1. *The vendor can strictly increase his profit by implementing a policy in which software pirates are restricted from patching security vulnerabilities when either*

(i) *the effective security risk is sufficiently high and piracy enforcement level is low, or*

(ii) *the consumer population's potential for piracy is low and the enforcement level for piracy is sufficiently high.*[9]

The results of Proposition 1 have interesting implications. Because users who remain unpatched are exposed to security risk, many security advocates argue that software vendors should provide pirates with access to security patches. Their reasoning is that if pirates are not allowed to patch in environments with high security risk, then the net value obtained by the legitimate users in the consumer population will be decreased. This risk will get reflected in the perceived value of the software for purchasers and will reduce their willingness to pay. As a result, a policy that does not let pirates patch can affect the vendor's profits negatively as well. Although this argument has an appeal, it is countered by the fact that, by not letting the pirates patch, the vendor may induce some consumers with high tendency for piracy to actually purchase the product, thereby generating revenue. Part (i) of Proposition 1 states that the vendor is strictly better off with a restricted patch policy when the enforcement level is low and the security risk is sufficiently high. When the expected loss incurred by pirating the product is low, a large number of Type H consumers choose to pirate the product. In this case, the vendor can price the software optimally

[8] A technical statement of this proposition is given in the proof of the proposition in the appendix, as it is for all propositions in this paper henceforth.

[9] Throughout the paper, if it is indicated that a result holds when a parameter is *sufficiently high* or *sufficiently large*, it is meant that there exists a bound such that if the parameter's value lies above that bound, then the result is true. The symmetric interpretation applies for the phrases *sufficiently low* and *sufficiently small*.

at a level that would convince some Type H consumers to purchase the software under the policy $\rho = nl$. This combination of policy and pricing allows the vendor to capture some revenue otherwise lost to piracy and can increase vendor profits, although it reduces the overall security of the product. However, as the proposition states, the vendor's gains with optimal pricing under a restrictive policy is higher than his gains from reduced negative network externalities on Type L consumers under his optimal pricing for an unrestrictive policy.

When the market potential for piracy is low and the level of enforcement is high (i.e., in an environment where one should expect the least amount of piracy activity), surprisingly, it is also optimal for the vendor to not let pirates patch as part (ii) of Proposition 1 states. When tendency for piracy in the consumer population is low (i.e., when the percentage of Type H consumers in the population is low), the vendor weighs heavily the effect of his pricing on Type L consumers. In this case, deciding on which patch policy to employ, the vendor must consider two effects from Type H consumers. First, under policy nl, if the vendor prices low enough (while maintaining $p > \pi_d c_d$), some Type H consumers who would become pirates may choose to purchase the product to avoid security risks, which can bring additional revenue to the vendor. On the other hand, a policy of letting pirates patch would increase patching in the overall population and make the product more secure. In turn, the value of the software product would increase, allowing the vendor charge a higher price. In comparison, a restricted patch policy would also require the vendor to decrease his price to a certain extent to coax the relatively higher-valuation Type H consumers to purchase the product, and this reduction in price may cause a decline in revenue from Type L consumers. Part (ii) of Proposition 1 states that, when the population's inherent tendency to pirate is low and the enforcement level is sufficiently high, the vendor can successfully implement a restricted patch policy by setting a sufficiently large price to shift to a market structure where some Type H customers with high valuations are now purchasing and patching. Furthermore, such a policy would result in higher profits compared to the policy where he lets the pirates patch, despite the optimal price under the latter policy being higher.

Before we end this section, we will discuss another case where a restrictive patch policy is optimal, namely, the case where the market's piracy tendency is high. Intuitively, as the potential for piracy becomes large, the user population is characterized by mostly Type H consumers. One factor that impacts the vendor's decision is the negative network effects stemming from unpatched usage by pirates on the purchasing population, which is minimized under policy l. As ν becomes large, under policy l, the vendor's gains in revenues from an increased purchasing population through the reduction of negative network externalities cannot exceed the revenues gained from convincing high-valuation Type H consumers to purchase the product under policy nl. Therefore, when the piracy tendency in the market is high, the vendor's optimal decision tends to be $\rho^* = nl$.[10] This argument is consistent with industry observations. In January 2005 Microsoft announced that participation in the normally optional pilot of its Windows Genuine Advantage (WGA) program became mandatory for Norwegian, Czech, and Chinese language versions, which are primarily used in the three respective countries known to have high piracy rates of Windows XP. This meant that users of these three versions were required to validate legitimacy of their licenses to receive security updates (Microsoft 2005a).[11] Thus, Microsoft's selective participation requirement is similar to effectively implementing a less restrictive patching policy for countries with lower piracy rates and a more restrictive one for countries with high piracy rates. The converse of the statement, however, is subtle: The optimal strategy for countries with relatively low tendency for piracy is complicated and hinges on the enforcement level of piracy in the market, as we will see when we contrast the results of this section and those of §4.3.

[10] A technical derivation of this argument is available from the authors upon request.

[11] In addition to the patching ban on software pirates, Microsoft also offered other anti-piracy programs in these countries, such as rebates to the users who were tricked into buying pirated copies while trying to buy legitimate ones (Evers 2005). Such programs are parts of additional measures Microsoft employs to combat piracy.

4.3. When Should a Vendor Not Restrict Security Patches for Pirates?

In §4.2 we have seen the market conditions under which a policy that restricts unlicensed copies of a software product is optimal. In this section we explore the conditions and market characteristics under which an unrestricted patch release policy (i.e., one that allows pirates to patch) would be optimal. The following proposition states the result.

Proposition 2. *The vendor can strictly increase his profit by implementing a policy in which software pirates are allowed to patch security vulnerabilities when the effective security risk is not too high or too low, and either*

(i) *the population's tendency for piracy and the piracy enforcement level are low, or*

(ii) *the patching costs are low and the enforcement level is not too high.*

Part (i) of Proposition 2 contrasts with part (ii) of Proposition 1 and demonstrates the importance of piracy enforcement level on the policy decision when the population's tendency for piracy is low. When ν is low, the vendor is again strongly concerned about the effect of policy on Type L users. However, unlike the case with a high level of enforcement ($\pi_d c_d$), when the enforcement level is low the vendor finds it optimal to increase patching among the population by allowing pirates to patch. In this case, in contrast to the situation with a high enforcement level where the vendor could use a substantially lowered price to reap the benefits of a patch restriction policy, it is more profitable for him to increase security in the face of an increased number of Type H consumers pirating. Lowering his price can bring revenues from Type H consumers but may dramatically reduce the revenues he receives from Type L consumers. When the enforcement level is low and $\rho = nl$, reducing his price to induce higher-value Type H consumers to buy the product is not as profitable as maintaining a market structure where no Type H consumer is a buyer and all pirating users are unpatched. The latter market structure, however, is dominated in profitability by the market structure induced with optimal pricing under $\rho = l$, because allowing pirates to patch increases the total mass of the patching population, consequently increasing the value of the product for consumers.

These results help explain Microsoft's current "let the pirates patch" approach in countries such as the United States where the piracy tendency is low (BSA-IDC 2005). Currently, in most developed countries, although the penalties when detected are high, the probability of being detected for many individual users (such as students and home users) is usually considerably low. Therefore, the effective enforcement level ($\pi_d c_d$) is low. As a result, instead of applying a restrictive policy that will not be effective at boosting revenues in such an environment (as we discussed above), Microsoft can choose to apply a loose patch release policy to minimize the negative network security externalities that consumers who are more likely to pirate the software impose on paying consumers. This policy also allows Microsoft to keep its software price high.

Part (ii) of Proposition 2 demonstrates the role of patching costs on the patch restriction policy decision. Patch management is generally a costly endeavor that uses a great deal of time and other resources. However, patching costs for vulnerabilities, though often being substantial, vary across vulnerability types, software applications, and vendors. In addition, a substantial amount of industry efforts is focused on reducing the patching costs, as evidenced by trends toward reducing deployment costs of security patches and corporate investment in patch management systems to better control the efficiency and cost of the entire security patch life cycle. Therefore, although the patching costs can never be zero, it is important to explore the implications of low patching costs on the patch release restriction policy for the vendor. When patching costs are small, the patching activity increases for both Type L and Type H users, yet not all users necessarily patch. Still, as a result of increased patching, network security and the value of the software product increase for users. Consequently, the vendor can charge a high price for the product under both patch release policies, $\rho = l$ and $\rho = nl$. Under a restricted patch release policy, however, given that the network environment is secure because of high patching activity, to convert the Type H consumers with relatively large valuations into purchasers, the vendor has to reduce his price significantly. This price reduction hurts his profit under the market structure in which higher-valuation Type H consumers are

induced to become purchasers, and, as a result, under policy $\rho = nl$, the vendor's optimal pricing induces a market structure where no Type H consumers buy or patch. On the other hand, under policy $\rho = l$, some Type H consumers are patchers, and the security of the network is higher. This increased security allows the vendor to capture a larger population for the same price relative to what he could achieve under policy $\rho = nl$. Thus, when patching costs are low, it is optimal for the vendor to not restrict unlicensed users from applying security patches.

5. Implications of Patch Restrictions on Network Security, Piracy Enforcement, and Welfare

As we established in §4, the vendor's optimal selection of the security patch restriction policy can have a dramatic effect on the consumer market equilibrium and structure. This strong impact on the shaping of consumer market structure can have important implications on both the vendor's treatment and valuation of the characteristics of his own software as well as on social welfare. In this section we explore the implications of patch restriction policies on security, policy decisions, and welfare.

When we evaluate implications of patch restrictions, we will evaluate the induced social welfare. The social welfare (i.e., the sum of all parties' surpluses) in equilibrium is given by

$$W(p) \triangleq \int_{\mathcal{V}} E_\theta\left[(v - C(v, \theta, \sigma^*))1_{\{(v>v_s \wedge \theta=H)\vee(v>v_b\wedge\theta=L)\}}\right] dv. \tag{7}$$

5.1. Vendor Incentives for Investment in Software Security

We begin by exploring the effect of patch restrictions on the vendor's incentives to invest in software security. Specifically, we will show that, when the tendency for piracy in the population is sufficiently high, the vendor will prefer a less secure software product to a more secure one, despite the fact that a more secure product has a higher value for consumers, allowing him to charge a higher price for the product. Normally, given a choice, one would expect a vendor to prefer a more secure product when there is no cost difference. However, we show that even when investment in improving software security is costless, under certain conditions the vendor would prefer a lower-security product, and, as a consequence, under patch restrictions his incentives to invest in security are reduced significantly.

Proposition 3. *When the piracy enforcement level is low and the piracy tendency in the population is sufficiently high, the profits for the vendor can be higher with a high-security-risk software product compared to a low-security-risk product even when improving the security of the product is costless.*

Figure 4 illustrates the result of Proposition 3 as well as its welfare implications. Confronted with a choice between making his product a low-effective-security-risk one (i.e., low $\pi_a\alpha$) by reducing its vulnerabilities and having a high-effective-security-risk product (i.e., high $\pi_a\alpha$), the vendor's incentives to improve the security of the product can be severely misaligned with the interests of social welfare. When the population's tendency to pirate is high, with a high-security-risk product, as we have discussed in §4.2 and as can be seen from the figure, the vendor prefers to have a restricted patch release policy because a high effective security risk increases the

Figure 4 Security Risk Used as a Strategic Tool by a Profit-Maximizing Vendor Who Enforces Security Patch Restrictions

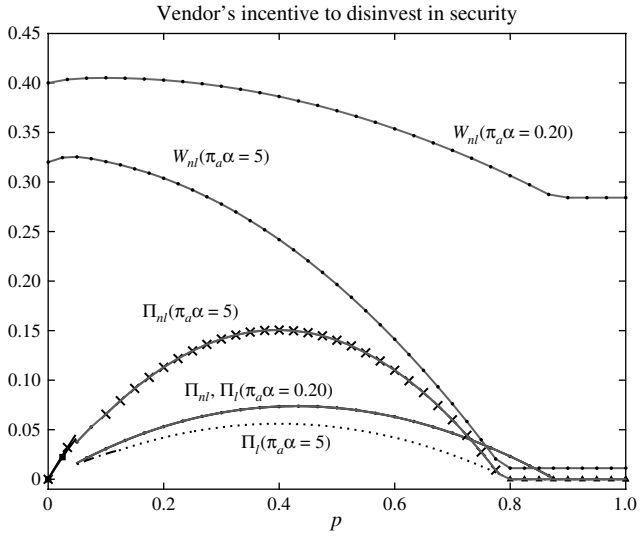

Notes. Profit and welfare curves are illustrated for low effective security risk ($\pi_a\alpha = 0.20$) and high effective security risk ($\pi_a\alpha = 5$) under both nonrestrictive ($\rho = l$) and restrictive ($\rho = nl$) security patch policies. The remaining parameters are $c_p = 0.20$, $\pi_d c_d = 0.05$, and $\nu = 0.65$.

incentives for conversion of Type H consumers into purchasers, which is profitable to the vendor. On the other hand, if the product is more secure, then, at the optimal price that the vendor offers the product, there will be decreased incentives for higher-valuation Type H consumers to purchase the product whether the vendor has a restricted or an unrestricted patching policy. As Proposition 3 states, this can result in lower vendor profits compared to the case where the effective security risk for the product is high and the vendor employs a restricted patch release policy (i.e., $\rho = nl$). Hence, the vendor is better off with a less secure product.

The vendor can affect the security risk by investing in efforts toward debugging security vulnerabilities during or after development. Many of the security after vulnerabilities that actually get exploited result from insufficient testing and debugging prior to the release of a product. Naturally, such security improvement actions are costly and time consuming. Proposition 3 states that, even if improving the security of the software were costless and even though his revenues suffer from the reduced security of his product, the vendor may find it preferable to not have a product with fewer vulnerabilities.[12] Rather, the vendor can use the lack of security in a product coupled with restricted security patch release policies as a strategic tool to increase profits. That is, a patch restriction policy can serve as a *substitute* to investing in the improvement of the security of his product. Such vendor behavior is detrimental to social welfare, as can also be seen in Figure 4, because increased risk engenders decreased usage, even among consumers who previously were purchasing the product.

Proposition 3 demonstrates an additional way security patch restrictions can hurt social welfare beyond the direct negative network externality effects. As a second and very important layer of negative effects on welfare, in certain cases, employment of patch restrictions reduces incentives to produce more secure software products. This is an important issue because the reduced security of a widely used software product (such as Windows XP) concerns and hurts the entire population. Thus, the impact and induced policy implications go well beyond vendor profits and merit serious attention by policy makers. This result also provides insights into both the observation that in recent years software vendors such as Microsoft tend to impose stricter security patch release policies in markets where piracy is a problem and the policy implications that result from such behavior (Microsoft 2006). Proposition 3 suggests that such restrictive policies can significantly reduce software security and hence bear a significant reduction in social welfare. In many cases, especially in emerging markets and for imported software, it is often perceived that piracy contributes to social welfare because the users can reap the benefits of the software without paying foreign companies for each copy used. However, for network software with potential vulnerabilities, our results establish that, in addition to legislating against restricted patch release policies, governments can find it desirable to increase piracy enforcement even though increasing enforcement reduces social surplus by decreasing total (legal or illegal) usage. That is, piracy enforcement has an increased benefit for governments as a remedy to increase network security in the face of potentially severe misalignment of incentives resulting from patch restriction policies especially in countries and markets where consumer tendency for piracy is high.

5.2. The Role of Piracy Enforcement

We next study the effect of piracy and patch restriction policies on the role of and policy determination for piracy enforcement. When considering the effect of the piracy enforcement level on vendor profits, one would think that an increased piracy enforcement level would help the vendor increase his profits. However, in the presence of negative network security effects and piracy in addition to the subsequent possibility of patch restriction policies, the effect of an increased piracy enforcement level is complex, and, surprisingly, an increased piracy enforcement level can even decrease vendor profits. The result is given by the following proposition.

Proposition 4. *Suppose that the effective security risk is high and the cost of patching is sufficiently small.*

[12] Note that, although the effective security risk $\pi_a \alpha$ is taken as a constant for the two levels (high and low) that Proposition 3 examines, the proposition still implies that the vendor would prefer not to invest in a more secure product, because it shows that switching to the more secure product, even if costless, would reduce the vendor's profit compared to that with the less secure product.

(i) *For low levels of piracy enforcement, increasing the piracy enforcement level can decrease vendor profits.*

(ii) *For high levels of piracy enforcement, increasing the piracy enforcement level increases vendor profits.*

There are important trade-offs that determine the impact of an increased piracy enforcement level on vendor profits. On one hand, an increased piracy enforcement level deters potential pirates, which has three positive effects on vendor profits. First, as a direct effect, an increased piracy enforcement level gives more flexibility to the vendor to increase his price without increasing piracy. Second, under a restrictive patch policy, a higher enforcement level increases the incentives for potential pirates to convert into purchasers and increases the vendor's profits. Third, stricter enforcement also deters some consumers who would pirate the product and choose not to patch from becoming users, thus exerting negative security externalities on other users. By deterring these consumers from usage, the vendor benefits from increased demand from legitimate users associated with a more secure product. On the other hand, with the additional presence of network security externalities, an increase in the network security level also has a downside for vendor profits. Specifically, for the case where the optimal vendor policy is employment of patch restrictions, increased network security reduces the incentives for users to convert from being pirates to being purchasers. The aggregate effect of an increased piracy enforcement rate depends on the joint pricing and patch restriction policy decision that the vendor makes in the face of these trade-offs. Figure 5 illustrates the effect of an increased piracy enforcement level on vendor profits resulting from the interplay of these factors and trade-offs. As can be seen from panel (a), for low software security risk levels, increased piracy enforcement helps increase vendor profits as one would normally expect.

However, for high security risk levels, the effect is non-monotonic as can be seen in panel (b). For low piracy enforcement levels, the market structure induced by the vendor's optimal pricing and policy combination implies that an increase in piracy enforcement and consequent reduction in network security risk lead to a loss in revenues from users who are no longer converted from pirates into purchasers. As a result, for low levels of enforcement, the vendor is forced to increase his price to recuperate lost revenue. However, this increase in price also hurts his revenues from users with low piracy tendencies. The aggregate effect is a decrease in his total return as stated in

Figure 5 The Impact of Piracy Enforcement on a Software Vendor's Price and Policy-Setting Behavior

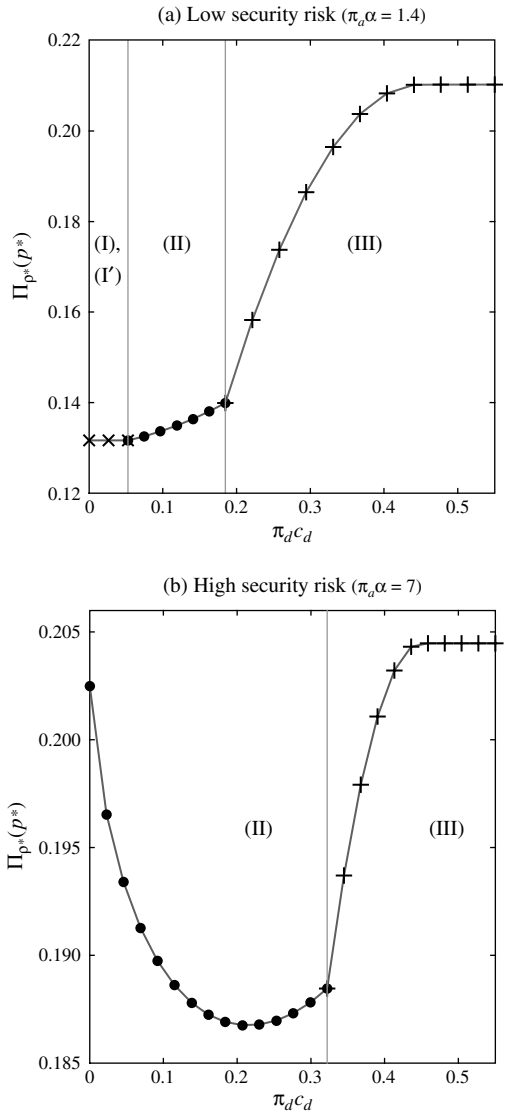

Notes. Optimal vendor profit curves are plotted as functions of $\pi_d c_d$ and exhibit varying returns characteristics dependent on market conditions. The parameters are $\pi_a\alpha = 1.4$, $c_p = 0.10$, and $\nu = 0.35$ for panel (a) and $\pi_a\alpha = 7$, $c_p = 0.10$, and $\nu = 0.35$ for panel (b). In the regions labeled (I) and (I′), the vendor is indifferent between security patch policies, and some consumers are pirating. Furthermore, the market structures are characterized by $0 < v_s < v_{sp} < v_b = v_p < 1$ and $0 < v_s < v_b = v_p < v_{sp} = 1$, respectively. In region (II), $\rho^* = nl$ and the market structure is characterized by $0 < v_s < v_b = v_p < v_{sp} < 1$. In region (III), the vendor is indifferent between policies and sets price to deter piracy.

part (i) of Proposition 4. In contrast, for higher enforcement levels and under the vendor's optimal pricing, conversion of consumers with higher valuations and high piracy tendencies, together with reduced network security externalities associated with increased enforcement levels, become dominant. As a result, vendor profits increase with increased enforcement level as stated in part (ii) of the proposition and as can also be seen in panel (b) of Figure 5.

The combination of patch restriction policies and the effect of negative security externalities in a network environment creates interesting welfare and policy implications. In particular, by inducing the vendor to strategically adjust his pricing to respond to changes in network security, increasing piracy enforcement can increase social welfare. The following proposition states the result.

Proposition 5. *When the effective security risk is sufficiently high, social welfare can continuously increase with increased piracy enforcement, in addition to jump increases at certain thresholds.*

Figure 6 demonstrates the effect of piracy enforcement on welfare. A common effect of increasing the enforcement level is the reduction of social welfare. This is because increased piracy enforcement reduces usage by consumers who pirate the product, and consequently the value generated for the economy by their usage is lost. In the presence of network security externalities, however, the effects become more complex. In particular, as Proposition 5 states, increasing piracy enforcement can surprisingly increase social welfare. As we discussed above, beyond a certain enforcement level, under a restrictive patch policy, increased enforcement reduces the unpatched pirated system population. In turn, the vendor finds it optimal to decrease the software price to entice an increased population of higher-valuation users with high piracy tendency to convert into purchasers. This increases overall usage and improves user patching behavior in certain population segments as well. As a consequence, increased piracy enforcement benefits social welfare as can be seen in both panels (a) and (b) of Figure 6. As depicted, increasing the enforcement level can also cause jump increases in welfare at certain critical levels by inducing the vendor to dramatically reduce his price to virtually eliminate piracy.[13] The possibility of a positive effect on welfare from increasing piracy enforcement level

Figure 6 The Impact of Piracy Enforcement on Social Welfare

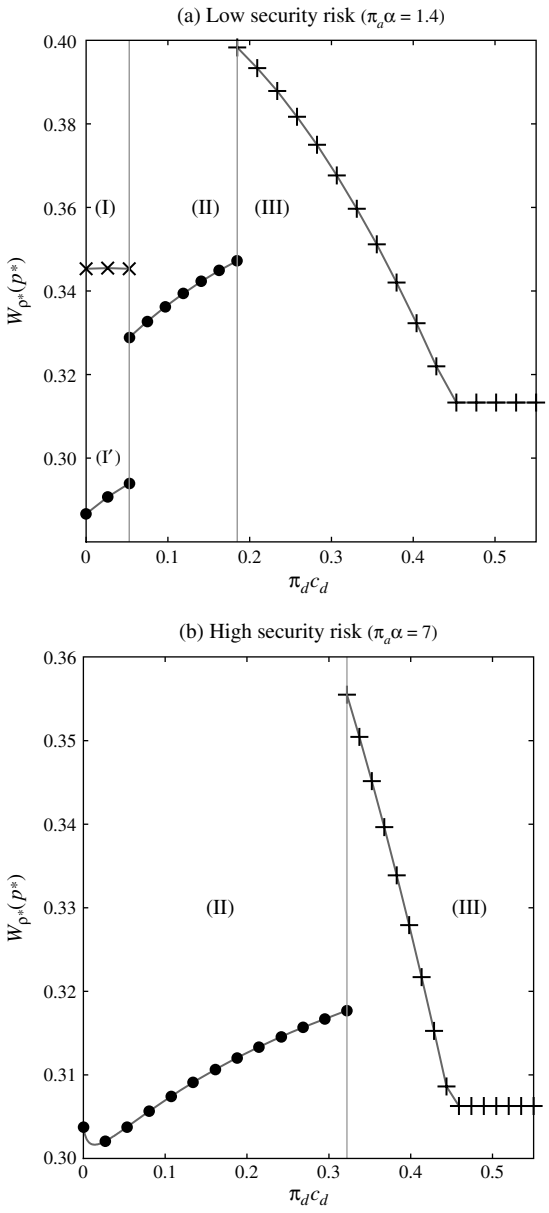

Notes. Welfare curves under a vendor's optimal policy and price setting are plotted as functions of $\pi_d c_d$. The regions marked (I) and (I′) denote welfare under policy $\rho = I$ and $\rho = nI$, respectively. The parameters and other regions are as described in the caption of Figure 5 for the corresponding panels.

[13] The shapes of the welfare curves in panels (a) and (b) of Figure 6 are robust over a significant range of the parameter space. For instance, a sensitivity analysis for panel (a) establishes that the jump in welfare between regions (II) and (III) remains for

has interesting policy implications. Governments are often reluctant to clamp down on piracy too forcefully because this would hurt the economy. However, under high security risk and network security externalities, in certain cases, increasing piracy enforcement can help increase welfare as well as vendor profits. Proposition 5 states that such increases can be desirable for governments, because they would not only better protect the intellectual property of software vendors but would also improve social surplus.

On the other hand, and interestingly, the story can be reversed for low piracy detection levels. For such a case, under high security risk, increasing piracy enforcement induces the vendor to increase his price. Hence, consumer usage decreases and welfare suffers, as can be seen in panel (b) of Figure 6. Combining this observation with the profit decrease on the same range (panel (b) of Figure 5) yields a policy implication that favors relaxed enforcement for low detection levels: When piracy enforcement is difficult, resulting in an effective enforcement level that is very low, it may be better to not increase the enforcement level for both vendor profits and social welfare.

Another interesting policy implication can be seen by examining panel (a) of Figure 5 wherein for low security risk levels the vendor is indifferent between a restrictive and an unrestrictive patch policy. Indifference arises naturally for such cases because, under both policies and the vendor's optimal pricing, the resulting market structures imply an absence of low piracy tendency users who do not patch and patching restrictions are not successful at converting any users with high piracy tendency into legitimate purchasers who patch. As a result, the vendor's patch restriction policy does not affect his profits. However, as panel (a) of Figure 6 demonstrates, the two policies have significantly different outcomes from the welfare perspective. Characteristically, a restrictive policy (i.e., $\rho = nl$) results in a substantial population of users with high piracy tendency who use the software and do not patch it. On the other hand, a permissive policy (i.e., $\rho = l$) induces a large segment of users with high piracy tendency to patch, thereby increasing both network security and the usage of the software. Therefore, the welfare under the permissive policy is higher than that under the restrictive policy. For such cases it would be advisable for policy makers to promote a permissive policy, because tipping the outcome toward increased usage and security not only raises the overall welfare but also makes all parties involved at least as well off, including the vendor. The promotion of a particular patching regime by policy makers to maximize the social value generated by software in an insecure network environment is an important subject that involves significant tradeoffs. We will examine this notion further in the next section.

5.3. Can Patch Restrictions Increase Social Welfare?

A large and important part of the debate on network software security patch restrictions focuses on the negative welfare implications of the vendor's restriction of security patches to only licensed users. The proponents of unrestricted patching policies argue that, when a vendor such as Microsoft does not allow owners of unlicensed copies of a software product access to security patches, the security of the entire network is reduced for all users. Therefore, the vendor's patch restrictions, even if they improve his profits, will have a negative effect on consumer surplus and social welfare. Hence, many in the IT community argue that the vendor should "do the right thing" in the social sense and allow software pirates to apply security patches (Schneier 2004). This is an appealing argument, and, as we have seen in §5.2, imposing of such a policy by a governing body can even be needed in certain cases to maximize social welfare. But does a policy of allowing pirates to patch *always* result in a higher consumer surplus let alone higher welfare? In this section we will demonstrate that, although in certain cases letting the pirates patch can result in an increase in social welfare, there are also cases where restrictive patch release policies can improve social welfare as well as consumer surplus. The following proposition summarizes our results.

Proposition 6. *When the enforcement level and the piracy tendency are low, social welfare is higher with an unrestricted patching policy. However, when the piracy*

$0<\pi_a\alpha<\infty$, $0<c_p<0.61$, and $0.04<\nu<1$; the slope discontinuity near $\pi_d c_d = 0.46$ persists for $0<\pi_a\alpha<\infty$, $0<c_p<1$, and $0<\nu<1$; and the jump in welfare between regions (I) and (II) occurs for $0.99<\pi_a\alpha<1.43$, $0.082<c_p<0.38$, and $0.25<\nu<0.36$.

enforcement level is sufficiently high, a patch release policy that restricts pirates from patching can increase welfare compared to an unrestricted policy.

Proposition 6 first states that, when the enforcement level is sufficiently low, an unrestricted patch release policy can increase social welfare. In such an environment, the vendor's optimal price under the restricted patch policy (i.e., $\rho = nl$) does not differ much from that under the policy where patching is unrestricted (i.e., $\rho = l$). Therefore, the social gains generated by the reduction of negative network security externalities by having increased patching in the Type H consumer population under policy $\rho = l$ exceed the loss of welfare from reduced consumer usage that results from a higher vendor optimal software price compared to the price under policy $\rho = nl$.

On the other hand, Proposition 6 also states that arguments that claim that restrictive patch release policies reduce social welfare are not necessarily true. In particular, when the piracy enforcement level ($\pi_d c_d$) is not too low, a restrictive patch policy can increase not only the social welfare but, more strongly, even the consumer surplus by reducing the vendor's price significantly. Figure 7 illustrates this result. When both the piracy enforcement level and patching costs are not too high, the vendor can profitably coax higher-valuation Type H consumers into buying under a policy that restricts pirates from patching because the network security risk that a pirate must assume by not being a licensed user can exceed the cost of being a patching purchaser. Hence, under such conditions, the vendor's optimal price can be in a region where there are Type H consumers who decide to purchase the product. Furthermore, weaker piracy enforcement can pull this price to a level that is significantly lower than the optimal price under the unrestricted patch permission policy. As a result, although the network security decreases with the policy $\rho = nl$ relative to the policy $\rho = l$, the generated additional surplus resulting from a lowered price can increase social welfare, as can also be seen in the figure. Note that, for the case demonstrated in the figure, consumer surplus also increases under the policy $\rho = nl$ because the vendor profit actually decreases with this policy compared to the policy $\rho = l$.

Figure 7 Increase in Social Welfare with the Imposition of a Restrictive Patch Release Policy

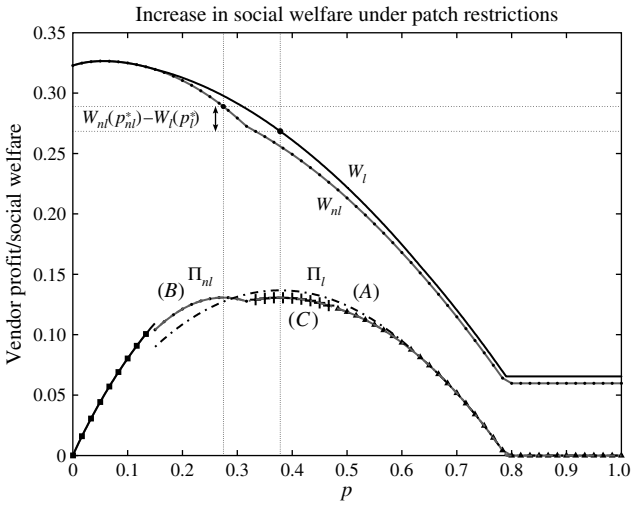

Notes. Profit and welfare curves are illustrated for both nonrestrictive ($\rho = l$) and restrictive ($\rho = nl$) security patch policies. The market structure denoted (A) is characterized by $0 < v_s < v_b < v_p = v_{sp} < 1$, the market structure denoted (B) is characterized by $0 < v_s < v_b < v_p < v_{sp} < 1$, and the market structure denoted (C) is characterized by $0 < v_s < v_b < v_p < v_{sp} = 1$, where in the latter two market structures v_{sp} denotes the threshold above which type H consumers purchase the product and patch. The parameters are $c_p = 0.21$, $\pi_a \alpha = 2.15$, $\pi_d c_d = 0.15$, and $\nu = 0.18$.

6. Concluding Remarks

In this paper we presented a model that allows us to simultaneously analyze the economics of network software security and software piracy. Specifically, we studied a currently debated important policy decision for software vendors and social planners, namely, whether to allow the owners of unlicensed copies of a software to apply security patches or not. Patch restriction policies not only affect the vendor's profits but also have a significant impact on social welfare through their implications on the vendor's pricing and policy decisions. Hence, the economics of determining the optimal policy is complex and must carefully take into account both software product characteristics and consumer market conditions. We identified the optimal patch restriction policies for the vendor and explored the effects and policy implications on software security, piracy enforcement, and social welfare. Our results showed that the joint effect resulting from the interaction of software piracy and network security externalities goes well beyond the separate effects of the two factors.

The joint consideration of these factors yields significant insights into economics of software policy and management.

One related economic issue is the utilization of standard quality differentiation methods to curb piracy (see, e.g., Alvisi et al. 2002, Sundararajan 2004). Although restriction of patching to legitimate users serves as a potential quality differentiator between purchased and pirated products, there are significant differences between our model and standard quality differentiation models. First, in standard quality differentiation models, each users' valuation is independent of other users' actions. Therefore, each user makes her decision independently. In contrast, in our case, because there are negative network externalities, each user's action depends on other users' actions, which must be accounted for in assessing each user's decision. As a result, unlike standard quality differentiation models, an equilibrium based on interdependence between cutoff valuations arises and needs to be solved. The solution of this equilibrium is central to the analysis. Second, the nature of the quality difference between alternatives from the consumer behavior perspective is different compared to standard quality differentiation models. Particularly, in our model, patching, which is the potential source of "quality differentiation," is a decision separate from purchasing. Consequently, even purchasers can be non-patchers and hence can choose *using* the "inferior quality" product even though they are purchasers. Third, in standard quality differentiation models, the vendor employs price differentiation to induce different types of customers to purchase different types of products. In our case there is only one price for the software product, and, even without this lack of price differentiation, different types of customers (e.g., high and low tendency for piracy) show very different purchase and patch behavior in equilibrium. In summary, our setting has fundamental differences compared to standard quality differentiation models in its nature, structure, economic drivers, and analysis. An interesting future study may combine different aspects of the two issues to explore the interaction of these approaches.

Another point to mention is that, like any other piece of software, security patches are also susceptible to hacking. This fact elicits the question whether consumers can use pirated versions of security patches to maintain the security of their systems as they can use pirated copies of the software itself. The key here is that there are differences between illegitimately obtaining and installing a base software product and illegitimately obtaining security patches. First, a user's pirating of the original product is merely a one-time effort. However, new security patches are deployed frequently and in large numbers continuously (Evers 2006, Microsoft 2007). Therefore, obtaining pirated security updates continuously, or pirated versions of a checking program that gets continuously updated, would require a much higher, ongoing effort level on a user's part. Second, software security patches are much more time critical than the software itself because if they are not downloaded and deployed in a given time window, the user may suffer an attack. Thus, it is not a good software maintenance strategy to wait for, search, and overall rely on pirated versions of security updates. Third, because of the continuous and repeated release of security updates, releasing such updates or their hacked versions over the web is the only viable means of effective provision. This forces users who want pirated versions of patches to rely on third-party websites and services that post these pirated/hacked security updates and that would inevitably be linked to hackers (Gantz et al. 2006). As one would expect, such websites are very unreliable and outright dangerous because they are likely to contain malware and intrusive software (see, e.g., Sovereign-Smith 2006). As a consequence, most users are well-advised to avoid this type of source for security patches and, understandably, do in fact stay away from them. In short, relying on pirated security updates for software maintenance is not likely to be a feasible or desirable strategy because it is costly and highly risky.

In our model the vendor's policy decision on whether to permit or restrict software pirates from accessing security patches is made in the first period. Although the patches are made available to certain users in the third period, the vendor's policy decision is credible and the outcome of the game is subgame perfect. This is because the vendor has no incentive to deviate from his patch release policy because his revenues are determined by the purchase behavior of

the users and not by the patching behavior. Another point to note is that we used a uniform demand distribution in our analysis. This distributional assumption is needed to maintain the tractability of the model. Furthermore, as is the case for most economic modeling, the goal of our study is to demonstrate economic arguments and derive insights from the analysis of those arguments. Our analysis with uniform demand achieves this goal.

One potential strategy that can improve the security of the software product and the network is the reduction of patching costs. However, reducing patching costs requires costly investments in software development, maintenance, and support. Therefore, a vendor would undertake reduction of patching costs as long as the returns justify the costs of investment. A potential future research study could explore the vendor's incentives in this area. A related direction for future research could also introduce costs of improving software security. However, note that our result in Proposition 3 shows that, even with zero security investment costs, the vendor can choose to have a less secure product over a more secure one. Another potential research topic could explore the effects of different types of security attacks. One example is the Distributed Denial of Service (DDoS) attack. In certain cases, unpatched systems can be taken over by hackers and coordinated to attack a given singled-out system or a website to prevent accessibility to users. These attacks can create a certain amount of congestion on the Internet, which can affect all users. However, such cases of congestion are usually short-lived and tend to directly affect only the users of a particular targeted site (Naraine 2002, Vijayan 2004). Overall, the general additional cost of congestion created is relatively small compared to the billions of dollars of costs associated with direct harm to users' systems in a widespread attack rather than a point-targeted attack, perhaps with the exception of the point target of the attack. Because of its separate economic structure and implications, such attacks (as well as other types of computer security attacks) deserve to be explored in separate studies. Such studies could be interesting avenues for future research.

In this paper we aimed to provide insights into the ongoing debate on whether a software vendor should restrict security patches only to legitimate users of a software product. Our results establish the criticality of the negative network security effects on policy decisions regarding patch and update management and the substantial effects that optimal selection of a network software security patch restriction policy can have on vendor profits and social welfare. Future research that continues to explore the network effects on software security can be valuable in shedding light on often controversial policy decisions concerning software security patch restrictions. Furthermore, such research may help to substantially improve the value generated by employing the right policy and limit the tremendous losses that piracy and software security vulnerabilities generate each year for the software industry and the economy in general.

Acknowledgments

We thank Sanjeev Dewan (the senior editor), the associate editor, and an anonymous referee as well as Haim Mendelson, Jim Patell, Jin Whang, and the seminar participants at University of California Berkeley, University of California at San Diego, Georgia Institute of Technology, Harvard Business School, the University of Maryland, the University of Rochester, Stanford University, and the Wharton School at the University of Pennsylvania for valuable comments and suggestions.

Appendix. Characterization of Consumer Market Equilibrium

Lemma A.1. *Suppose $\rho = l$ in which case there are no restrictions on accessing security patches for any user. Given $p \in (\pi_d c_d, 1]$, there exists a unique equilibrium in the consumer market. The equilibrium consumer strategy profile is characterized by four cutoff values $0 < v_s < v_b$, $v_{sp} \le v_p \le 1$ such that*

$$\sigma^*(v, L) = \begin{cases} (NU, NP) & \text{if } 0 \le v < v_b; \\ (B, NP) & \text{if } v_b \le v < v_p; \\ (B, P) & \text{if } v_p \le v \le 1, \end{cases} \tag{8}$$

and

$$\sigma^*(v, H) = \begin{cases} (NU, NP) & \text{if } 0 \le v < v_s; \\ (S, NP) & \text{if } v_s \le v < v_{sp}; \\ (S, P) & \text{if } v_{sp} \le v \le 1. \end{cases} \tag{9}$$

Proof. For Type L consumers, each consumer can take one of three actions: (NB, NP), (B, NP), or (B, P). If a consumer with valuation v chooses to not buy, then her payoff is zero. If she chooses to buy but not patch, and the total mass of unpatched population is u, her expected payoff is $v - p - \pi_a u \alpha v$. If she chooses to buy and patch, her total

expected payoff is $v - p - c_p$. Hence, a consumer prefers patched usage over unpatched usage if and only if her valuation satisfies $v \geq c_p/(\pi_a \alpha u)$. Therefore, if it exists, the segment of consumers who patch will be the one with highest valuations, and these two segments cannot be reversed. A consumer prefers unpatched usage to non-usage if and only if $v \geq p/(1-\pi_a \alpha u)$. Therefore, the segment of non-users will be the one with lowest valuations, and, if a segment of unpatched users exists, it will be the segment in between. Because $p > \pi_d c_d$, it is possible that there is either no segment of patched users or no segment on unpatched users. Therefore, the equilibrium consumer strategy profile for Type L consumers is characterized by $0 < v_b \leq v_p \leq 1$ such that for all $v \in \mathcal{V}$, $\sigma^*(v, L)$ satisfies (8). Because under different market structures some user segments may not exist, we utilize more general expressions to build up the characterization of v_b and v_p. Specifically, a Type L consumer buys the product and patches in the second period in case a security vulnerability is revealed if and only if

$$v \geq \max\left(\frac{c_p}{\pi_a u \alpha}, p + c_p\right) \tag{10}$$

and buys the software if and only if

$$v \geq \min\left(\frac{p}{1-\pi_a \alpha u}, p + c_p\right). \tag{11}$$

Now, note that the following three inequalities are algebraically equivalent

$$\begin{aligned} \frac{p}{1-\pi_a \alpha u} \geq p + c_p \quad &\Longleftrightarrow \quad \frac{p}{1-\pi_a \alpha u} \geq \frac{c_p}{\pi_a \alpha u} \\ &\Longleftrightarrow \quad p + c_p \geq \frac{c_p}{\pi_a \alpha u}, \end{aligned} \tag{12}$$

Because,

$$\begin{aligned} \frac{p}{1-\pi_a \alpha u} \geq p + c_p \quad &\Longleftrightarrow \quad p \pi_a \alpha u \geq c_p - c_p \pi_a \alpha u \\ &\Longleftrightarrow \quad \frac{p}{1-\pi_a \alpha u} \geq \frac{c_p}{\pi_a \alpha u}, \end{aligned} \tag{13}$$

and,

$$\begin{aligned} \frac{p}{1-\pi_a \alpha u} \geq \frac{c_p}{\pi_a \alpha u} \quad &\Longleftrightarrow \quad p \pi_a \alpha u + c_p \pi_a \alpha u \geq c_p \\ &\Longleftrightarrow \quad p + c_p \geq \frac{c_p}{\pi_a \alpha u}. \end{aligned} \tag{14}$$

Then, when $v_b < 1$ and $v_p < 1$, by (10) and (11), in equilibrium

$$v_p = \max\left(\frac{c_p}{\pi_a \alpha u(\sigma^*)}, p + c_p\right), \tag{15}$$

and

$$v_b = \min\left(\frac{p}{1-\pi_a \alpha u(\sigma^*)}, p + c_p\right). \tag{16}$$

Thus, by (12), (15), and (16), either both $v_b = p + c_p$ and $v_p = p + c_p$ or $v_b < v_p$, in which case

$$v_b = \frac{p}{1-\pi_a \alpha u(\sigma^*)}, \tag{17}$$

and

$$v_p = \frac{c_p}{\pi_a \alpha u(\sigma^*)}. \tag{18}$$

Because $p > \pi_d c_d$, any Type H consumer who decides to use the software is pirating rather than purchasing. If a Type H consumer with valuation v decides to pirate and patch the software product, her expected payoff is $v - \pi_d c_d - c_p$. If she decides to pirate but not patch, her expected payoff is $v - \pi_d c_d - \pi_a u \alpha v$. Therefore, she pirates the product and patches in the second period in case a security vulnerability is revealed if and only if

$$v \geq \max\left(\frac{c_p}{\pi_a u \alpha}, \pi_d c_d + c_p\right). \tag{19}$$

Consequently, if it exists, the segment of Type H users who patch will be the one with highest valuations. Hence, there exists a $v_{sp} \in (0, 1]$ such that a Type H consumer with valuation $v \in \mathcal{V}$ will pirate and patch if and only if $v \geq v_{sp}$, in which case $\sigma^*(v, H) = (S, P)$.

Consider the decision to pirate in the first period. If a Type H consumer with valuation v decides to pirate the product, she will incur an expected loss associated with piracy of $\pi_d c_d$. Thus, she will pirate the software if and only if

$$v \geq \min\left(\frac{\pi_d c_d}{1-\pi_a \alpha u}, \pi_d c_d + c_p\right), \tag{20}$$

and, therefore, all Type H consumers with valuations above a threshold value, $v_s \in (0, 1]$, will pirate the software.

By (11) and (20), $v_s < v_b$ because $p > \pi_d c_d$. By definition, $v_{sp} \geq v_s$. Suppose that $0 < v_s = v_{sp} < 1$ and $c_p > 0$. Further suppose that $v_b = v_p$. Then, there exists $0 < v < v_s$ such that $v \geq \pi_d c_d + C(v, H, \sigma^*) = \pi_d c_d$, which is a contradiction. Suppose that $v_p > v_b$. Then, $c_p > \pi_a \alpha u v_b$, which implies that a Type H consumer with valuation v_b should not be patching, which is a contradiction because $v_b > v_s = v_{sp}$. Therefore, we conclude that, when $\pi_d c_d < p < 1$, there exist $0 < v_s < v_{sp} \leq 1$ satisfying (9). Hence, in equilibrium, when $v_{sp} < 1$, by (19) and (20) we have

$$v_s = \frac{\pi_d c_d}{1-\pi_a \alpha u(\sigma^*)}, \tag{21}$$

and

$$v_{sp} = \frac{c_p}{\pi_a \alpha u(\sigma^*)}. \tag{22}$$

Finally, we will establish that $v_{sp} \leq v_p$. Suppose that $v_p < 1$. Furthermore, suppose that $v_p > v_b$. Then, by (12) and (15), v_p satisfies (18). A Type H consumer with valuation v_p faces the same patching trade-off; hence, $v_{sp} = v_p$. Suppose that $v_b = v_p$. Then, by (12) and (22), $v_p = p + c_p \geq v_{sp}$. This completes the proof. □

Lemma A.2. *Suppose $\rho = nl$ in which case software pirates are restricted from applying security patches for vulnerabilities. Given $p \in (\pi_d c_d, 1]$, there exists a unique equilibrium in the consumer market. The equilibrium consumer strategy profile is characterized by four cutoff values $0 < v_s < v_b \le v_p \le v_{sp} \le 1$ such that $\sigma^*(v, L)$ satisfies* (8) *and*

$$\sigma^*(v, H) = \begin{cases} (NU, NP) & \text{if } 0 \le v < v_s; \\ (S, NP) & \text{if } v_s \le v < v_{sp}; \\ (B, P) & \text{if } v_{sp} \le v \le 1. \end{cases} \tag{23}$$

Proof. The equilibrium behavior of Type L consumers is similar to that given in Lemma A.1. Under policy nl, Type H consumers can patch in case of a vulnerability only if they have legitimately purchased the software. Because $p > \pi_d c_d$, no Type H consumer will find it optimal to purchase and not patch. Therefore, the equilibrium strategy of a Type H consumer with valuation v who decides to use the software must be either (S, NP) or (B, P). If she chooses (B, P), then her expected payoff is $v - p - c_p$. If she chooses (S, NP), then her expected payoff is $v - \pi_d c_d - \pi_a \alpha u v$. Thus, her equilibrium strategy is (B, P) if and only if

$$v \ge \max\left(\frac{p + c_p - \pi_d c_d}{\pi_a \alpha u}, p + c_p\right). \tag{24}$$

Consequently, if a segment of Type H users who buy and patch exists, then this segment will be characterized by the users with highest valuations. Furthermore, its order relative to the segment of Type H users who pirate and do not patch cannot be reversed. Hence, there exists a $v_{sp} \in (0, 1]$ such that for all $v \in \mathcal{V}$, $\sigma^*(v, H) = (B, P)$ if and only if $v \ge v_{sp}$. On the other hand, a Type H consumer will either pirate or purchase the software if and only if

$$v \ge \min\left(\frac{\pi_d c_d}{1 - \pi_a \alpha u}, p + c_p\right). \tag{25}$$

Furthermore, all Type H consumers with valuations higher than a certain $v_s \in (0, 1]$ will use the software.

Now suppose (21) does not hold. Then, by (25), $\pi_d c_d / (1 - \pi_a \alpha u) > p + c_p$ if and only if $p + c_p > (p + c_p - \pi_d c_d)/(\pi_a \alpha u)$, which implies that $v_s = v_{sp} = p + c_p$. Furthermore, because $p > \pi_d c_d$, by (12), $v_b = v_p = p + c_p$, which is a contradiction. Therefore, (21) is satisfied, and hence $v_s < v_b$. By definition, we have $v_{sp} \ge v_s$. Hence, when $\pi_d c_d < p < 1$ there exist $0 < v_s \le v_{sp} \le 1$ satisfying (23). Furthermore, because $\pi_d c_d / (1 - \pi_a \alpha u) \le p + c_p$, by (24), under policy nl, we must have

$$v_{sp} = \frac{p + c_p - \pi_d c_d}{\pi_a \alpha u(\sigma^*)}, \tag{26}$$

when $v_{sp} < 1$. Now suppose $v_{sp} < 1$. Then, by (15), (24), and (26), we have $v_p < 1$. If $v_b < v_p$, by (12) and (15), $v_p = c_p / (\pi_a \alpha u(\sigma^*))$, which implies $v_{sp} > v_p$ because $p > \pi_d c_d$. When $v_b = v_p$, again $v_p = p + c_p$, and by (24), $v_{sp} \ge v_p$. This completes the proof. □

References

Alvisi, M., E. Argentesi, E. Carbonara. 2002. Piracy and quality choice in monopolistic markets. German Working Papers in Law and Economics, University of Mannheim, Mannheim, Germany.

Anderson, R., T. Moore. 2006. The economics of information security. *Science* **314**(5799) 610–613.

Arora, A., J. P. Caulkins, R. Telang. 2006. Research note–sell first, fix later: Impact of patching on software quality. *Management Sci.* **52**(3) 465–471.

Arora, A., R. Telang, H. Xu. 2005. Optimal policy for software vulnerability disclosure. Working paper, Carnegie Mellon University, Pittsburgh.

August, T., T. I. Tunca. 2006. Network software security and user incentives. *Management Sci.* **52**(11) 1703–1720.

Bakos, Y., E. Brynjolfsson, D. Lichtman. 1999. Shared information goods. *J. Law Econom.* **XLII** 117–155.

Bass, D. 2005. Microsoft to begin blocking fake Windows users from downloads. *Bloomberg* (January 26), http://www.bloomberg.com.

Besen, S. M., S. N. Kirby. 1989. Private copying, appropriability, and optimal copying royalties. *J. Law Econom.* **32** 255–280.

Bloor, B. 2003. The patch problem: It's costing your business real dollars. *Baroudi Bloor*, http://www.mithras.itworld.com/download/special_reports/smallbusiness/PatchProblemReport_BaroudiBloor.pdf.

BSA-IDC. 2005. Second annual BSA and IDC global software piracy study. (May 25).

Cavusoglu, H., H. Cavusoglu, S. Raghunathan. 2004a. How should we disclose software vulnerabilities. (December) *Workshop on Information Technology and Systems* (*WITS*), Washington, D.C.

Cavusoglu, H., H. Cavusoglu, J. Zhang. 2004b. Security patch management: Can't live with it, can't live without it. (December) *Workshop on Information Technology and Systems* (*WITS*), Washington, D.C.

Cavusoglu, H., B. Mishra, S. Raghunathan. 2005. The value of intrusion detection systems in information technology security. *Inform. Systems Res.* **16**(1) 28–46.

Chen, Y.-N., I. Png. 2003. Information goods pricing and copyright enforcement: Welfare analysis. *Inform. Systems Res.* **14**(1) 107–123.

Connor, K. R., R. P. Rumelt. 1991. Software piracy: An analysis of protection strategies. *Management Sci.* **37**(2) 125–139.

Curien, N., G. Laffond, J. Laine, F. Moreau. 2004. Towards a new business model for the music industry: Accomodating piracy through ancillary products. *Laboratoire d'Econometrie Conservatoire National de Arts et Metiers* **1**(4–5) 1–21.

Evers, J. 2005. Microsoft to require Windows piracy check. *PC World* (January 26) http://www.pcworld.com/article/id,119458-page,1/article.html.

Evers, J. 2006. Another hefty patch month for Microsoft. *CNET News.com* (August 8).

Fried, I. 2005. Piracy-check mandatory for Windows add-ons. *CNET News.com* (July 25).

Gantz, J. F., A. Gillen, C. A. Christiansen. 2006. The risks of obtaining and using pirated software. (October) *Microsoft.com*.

Gopal, R. D., G. L. Sanders. 1997. Preventive and deterrent controls for software piracy. *J. Management Inform. Systems* **13**(4) 29–47.

Harbaugh, R., R. Khemka. 2001. Does copyright enforcement encourage piracy? Working Paper 2001-14, Claremont McKenna College, Claremont, CA.

Holsing, N. F., D. C. Yen. 1999. Software asset management: Analysis, development and implementation. *Inform. Resources Management J.* **12**(3) 14–26.

Hou, C. H. 2004. Security for all. *Computer Times* (May 5). http://computertimes.asia1.com.sg/news/story/0,5104,2292,00.html.

Jaisingh, J., Q. Li. 2005. The optimal time to disclose software vulnerability: Incentive and commitment. (November) Working paper, Hong Kong University of Science and Technology, Kowloon, Hong Kong.

Johnson, W. 1985. The economics of copying. *J. Political Econom.* **93**(1) 158–174.

Liebowitz, S. J. 1985. Copying and indirect appropriability: Photocopying of journals. *J. Political Econom.* **93**(5) 945–957.

Microsoft. 2002. Windows XP product activation. (August) *Microsoft.com*.

Microsoft. 2005a. Microsoft to implement worldwide anti-piracy initiative. (January) *Microsoft.com*.

Microsoft. 2005b. Windows genuine advantage 1.0 goes live. (July) *Microsoft.com*.

Microsoft. 2006. Founder signs genuine Windows cooperative engagement agreement with Microsoft (April) *Microsoft.com*.

Microsoft. 2007. Microsoft security bulletin summaries and webcasts. (January) *Microsoft.com*.

Moore, D., C. Shannon, J. Brown. 2002. Code-red: A case study on the spread and victims of an internet worm. *Proc. 2nd ACM SIGCOMM Workshop on Internet Measurement*, ACM Press, New York, 273–284.

Naraine, R. 2002. Massive DDoS attack hit DNS root servers. (October) *InternetNews.com* (October 23).

Novos, I. E., M. Waldman. 1984. The effects of increased copyright protection: An analytic approach. *J. Political Econom.* **92**(2) 236–246.

Peitz, M., P. Waelbroeck. 2003a. Making use of file sharing in music distribution. Mimeo, University of Mannheim, Mannheim, and ECARES, Germany, Free University of Brussels, Brussels, Belgium.

Peitz, M., P. Waelbroeck. 2003b. Piracy of digital products: A critical review of the economics literature. CESinfo Working Paper Series, No. 1071, Ludwig-Maximilians-Universitaet, and the Ifo Institute for Economic Research, Munich, Germany.

Rao, S. 2003. Copyright: Its implications for electronic information. *Online Inform. Rev.* **27**(4) 264–275.

Rooney, P. 2005. Channel praises Microsoft plan to "stigmatize" software pirates. (January 26) *CRN.com*.

Salloway, M. 2004. Common issues with Windows XP service pack 2. *Mark Salloway's Windows XP Resource Center*. http://www.mvps.org/marksxp/WindowsXP/SP2/common.php#sp2.

Schneier, B. 2004. Microsoft's actions speak louder than words. (May) *NetworkWorld.com*.

Shy, O., J.-F. Thisse. 1999. A strategic approach to software protection *J. Econom. Management Strategy* **8**(2) 163–190.

Slive, J., D. Bernhardt. 1998. Pirated for profit. *Canadian J. Econom.* **31**(4) 886–899.

Sovereign-Smith, R. 2006. The net threat. *Digit* (August) 108–110.

Sundararajan, A. 2004. Managing digital piracy: Pricing and protection. *Inform. Systems Res.* **15**(3) 287–308.

Takeyama, L. N. 1994. The welfare implications of unauthorized reproduction of intellectual property in the presence of demand network externalities. *J. Indust. Econom.* **62** 155–166.

Takeyama, L. N. 1997. The intertemporal consequences of unauthorized reproduction of intellectual property. *J. Law Econom.* **40**(2) 511–522.

Varian, H. 2000. Buying, renting and sharing information goods. *J. Indust. Econom.* **48** 473–488.

Varian, H. R. 2005. Copying and copyright. *J. Econom. Perspectives* **19**(2) 121–138.

Vijayan, J. 2004. Akamai now says it was targeted by DDoS attack. Computerworld (June) Computerworld.com.

Weaver, N., V. Paxson, S. Staniford, R. Cunningham. 2003. A taxonomy of computer worms. *Proc. 2003 ACM Workshop on Rapid Malcode*, ACM Press, New York, 11–18.

Worthington, D. 2004. Microsoft: SP2 will not install on pirated copies of XP. *BetaNews*. (May 11), http://www.betanews.com/article/1084264398.