

Let the Pirates Patch?

An Economic Analysis of Network Software Security Patch Restrictions

Terrence August and Tunay I. Tunca
Graduate School of Business, Stanford University
Stanford, CA, 94305

EXTENDED ABSTRACT

1 Motivation

Piracy has long been an important concern for the software industry. The relative ease of reproducing 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 U.S. alone is about \$7 Billion per year, and the annual global cost of piracy exceeds \$30 Billion (Rooney 2005). Further, 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 February 27, 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 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. 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., Wagner 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 five months after the first announcement in February, at the worldwide launch of the 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.

At the heart of the issue lies the difficult decision that not only Microsoft but also any software vendor faces. Prohibiting the owners 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. The result of this critical policy decision has pronounced economic consequences for the vendor.

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 on the optimal vendor patch release policy under software piracy in connection with current empirical observations and the ongoing debate. Building on the base model in August and Tunca (2006), we explore the economic 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. Specifically, our analysis has three main purposes. First, we identify the conditions under which each policy will be optimal for a software vendor. Second, we show that given the ability to use software security restrictions, a vendor may prefer a less secure product and hence can have reduced incentives to invest in improving software security. Third, 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.

2 Overview of the Model, Analysis, and Results

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) by introducing heterogeneity in consumer tendencies toward piracy. Specifically, each consumer is assigned a type

from $\Theta = \{L, H\}$, where Type L denotes a consumer of “Low piracy tendency” (or no tendency in our case for simplicity), while Type H denotes a consumer of “High piracy tendency”. For any given consumer, the probability that she is of Type H is given by $\nu \in [0, 1]$. There are four time periods. In the first period, the vendor sets a policy in regard to security patch restrictions for unlicensed users and sets the price 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, whether the software has a security vulnerability or not is revealed. 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 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 $c_p > 0$, 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).

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 $\pi_a u$. Further, if a user’s system goes unpatched and is hit by the attack, one would expect that she suffers a loss positively correlated with her valuation. That is, consumers with high valuations will suffer higher losses than consumers with lower valuations due to opportunity costs, higher criticality of data and loss of business. For simplicity, we assume that the correlation is of first order such that the loss that a consumer with valuation v suffers if she is hit is αv where $\alpha > 0$ is a constant.

We analytically demonstrate that when software is highly risky and the piracy enforcement level is low, or the population’s tendency for piracy is high, it is optimal for the vendor to impose security patch restrictions on unlicensed users. Although not permitting pirates to patch indeed reduces the product’s security and hence its perceived value by the users, we show that when the enforcement level is low, under high effective security risk, it is optimal for the vendor to restrict security patches only to licensed users. Similarly, when the probability of being a consumer with a tendency for piracy is sufficiently high, revenues generated by incentivizing potential pirates to purchase the software via patch restrictions exceed those stemming from a reduction in negative security network externalities brought about by an unrestricted patch release policy. However, the opposite is not necessarily true. 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, while in an environment with a low level of piracy enforcement, he should employ an unrestricted patching policy. The magnitude of the cost of patching also plays an important role. We show that when

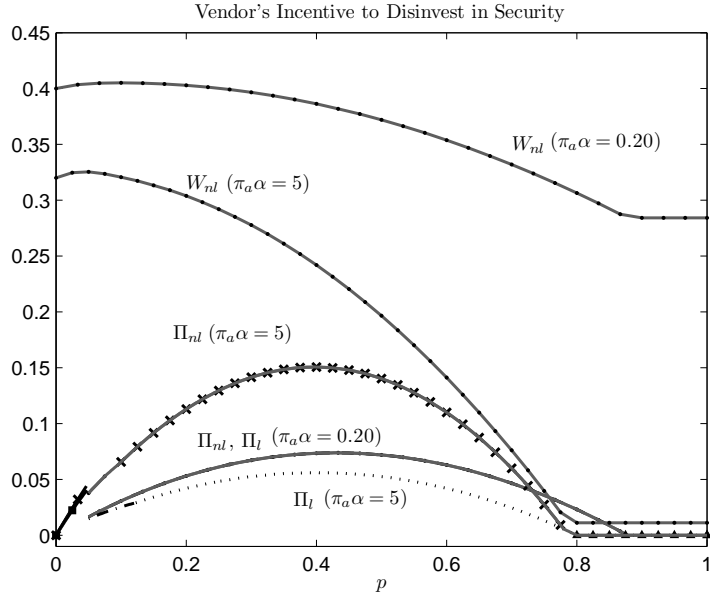


Figure 1: Security risk used as a strategic tool by a profit maximizing vendor who enforces security patch restrictions. Profit and welfare curves are illustrated for low effective security risk ($\pi_a \alpha = 0.20$) and high effective security risk ($\pi_a \alpha = 5$) under both non-restrictive ($\rho = l$) and restrictive ($\rho = nl$) security patch policies. Each pattern for each profit curve denotes a different consumer market structure. The remaining parameters are $c_p = 0.20$, $\pi_d c_d = 0.05$ and $\nu = 0.65$.

patching costs are sufficiently low, contrary to what one might think, the profit maximizing policy for the vendor is to allow all users, licensed or unlicensed, to apply security patches.

Next, we explore the effect of security patch restrictions on the vendor's incentives to improve the security of his product. We analytically establish that when the probability of each user being a potential pirate is not too low and the piracy enforcement level is low, then the vendor prefers to have a software product with *high security risk* and employ a restricted patch policy compared to having a product with low security risk as illustrated in Figure 1. Considering that this conclusion remains true even with no cost for security improvements, it implies that in environments where piracy laws are not well enforced, security patch restrictions can be used as a substitute to investment in software security improvements. As a result, we show that social welfare can suffer substantially which provides governments with a new reason to increase piracy enforcement levels.

Finally, we explore the optimality of patch restriction policies from a social welfare point of view. We find that when the piracy enforcement level is low, an unrestricted patching policy tends to be optimal for maximizing social welfare. However, contrary to some arguments given in the IT community, we show that restricting patching to only licensed users can actually improve social welfare. Such an increase can occur especially when the piracy enforcement level is sufficiently high. In such a case, the vendor may find it optimal to price the product low under a restricted patching policy in comparison to an unrestricted patching policy. Therefore, even though the security of the network may decrease, the total usage and social surplus generated by the software in the economy may go up, as seen in Figure 2. Additionally, the vendor might not choose this joint price/policy strategy when left to his own decision, thus it may be optimal for a social planner to explicitly impose patch restriction policies to improve social surplus in these instances. Our results demonstrate that software security patch restrictions can have important economic consequences

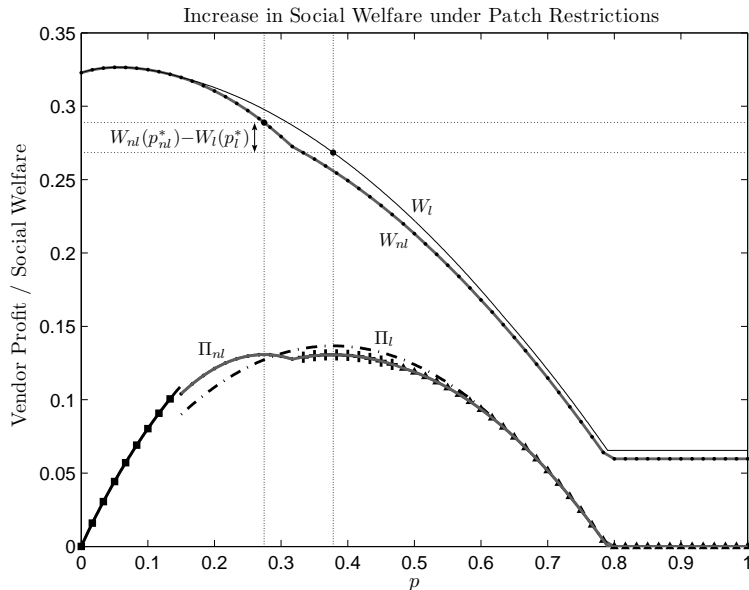


Figure 2: Increase in social welfare with the imposition of a restrictive patch release policy. Profit and welfare curves are illustrated for both non-restrictive ($\rho=l$) and restrictive ($\rho=nl$) security patch policies. Each pattern for each profit curve denotes a different consumer market structure. The parameters are $c_p=0.21$, $\pi_a\alpha=2.15$, $\pi_dc_d=0.15$, and $\nu=0.18$.

on the incentives of the vendor and the users and hence on vendor profits and welfare.

References

- August, T. and T. I. Tunca (2006). Network software security and user incentives. Forthcoming, *Management Science*.
- Bloor, B. (2003). The patch problem: It's costing your business real dollars. *Baroudi Bloor*. <http://www.baroudi.com/pdfs/patch.pdf>.
- BSA-IDC (2005, May 25). Second annual BSA and IDC global software piracy study. Available at BSA.org.
- Fried, I. (2005, Jul 25). Piracy-check mandatory for windows add-ons. CNET News.com.
- Moore, D., C. Shannon, and J. Brown (2002). Code-red: a case study on the spread and victims of an internet worm. *Proceedings of the Second ACM SIGCOMM Workshop on Internet Measurement*, 273–284.
- Rooney, P. (2005, Jan 26). Channel praises microsoft plan to 'stigmatize' software pirates. CRN.com.
- Schneier, B. (2004, May 31). Microsoft's actions speak louder than words. Network World. www.networkworld.com.
- Wagner, M. (2004, May 13). Even thieves need protection. www.securitypipeline.com.
- Weaver, N., V. Paxson, S. Staniford, and R. Cunningham (2003). A taxonomy of computer worms. *Proceedings of the 2003 ACM Workshop on Rapid Malcode*, 11–18.