

Competition and quality restoration: An empirical analysis of vendor response to software vulnerabilities

Ashish Arora, Chris Forman, Anand Nandkumar¹ and Rahul Telang

{ashish,cforman, anandn,rtelang}@andrew.cmu.edu

Carnegie Mellon University

1. Introduction

Costs related to information security have recently had a large and increasing impact on the U.S. economy. A recent study put the annual cost of major software bugs to the U.S. economy at over \$60 billion (NIST 2002). Though there are not as yet any official U.S. government statistics on information security, several private groups have demonstrated the growth in security-related incidents and their antecedents².

The rapid increase in the number of vulnerabilities discovered in software over the past several years has led many to argue that high levels of concentration and significant early mover advantages in software markets lead to an under-provision of software quality. Firms with market power could deliberately under provide quality in an effort to maximize profits. However, others have argued that the link between market structure and provision of software quality has been exaggerated. For one, some studies have found that users are unwilling to pay for software quality because it is difficult for them to value it ex ante: If under provisioning of quality is due to lack of user willingness to pay then market structure should have little impact on quality. Second, when software includes components that are common to several software markets, changes in the market structure of one market may influence patching behavior in another. Thus, patching behavior of software producers that operate in concentrated markets may be influenced by increases in competition in markets that share complementary components.

In this paper, we examine the relationship between competition and one dimension of software quality: the duration of patch release time to known vulnerabilities. We begin by developing our hypotheses in a model of vendor patching behavior that is based on vendors' internalization of end user losses and vendors' costs of developing patches. Increased market competition increases vendor losses from unpatched vulnerabilities. As a result, higher competition leads to lower patching times. We label this direct impact of competition as the *competition effect*. Further, increases in the total number vendors—both competitors and noncompetitors—that share the vulnerability will lead (in probability) to earlier initial public disclosure of the vulnerability. We label this relationship the *disclosure effect*.

We demonstrate that competition and disclosure each have an economically and statistically significant impact on patching times. We show that a 10% increase in the number of rivals lowers expected patching times between 2.7% to 5.4%; translating to a decline of 7 to 25 days due to the direct effects of competition. We also show that disclosure plays a role: a 10% increase in vendors from unrelated markets that share the same vulnerability will lead to a decrease in expected patching times of 0.9% to 2.1%. These elasticities translate into a decline of

¹ Corresponding author. We thank CERT/CC and Harte Hanks Market Intelligence for providing essential data.

² (1) The number of information security incidents reported to security research center CERT/CC, a federally funded research laboratory that researches Internet security problems, grew from 2412 in 1995 to 137,529 in 2003. (2) Meanwhile, the number of reported software vulnerabilities, one leading indicator of security incidents, grew from 171 in 1995 to 5990 in 2005.

12 to 15 days. Last, we also show that market size will also increase quality provision: a 10% increase in vendor installed base leads to a 1.3% to 1.4% decline in patching times.

Our research is unique in demonstrating how products with common technological inputs can influence output market competition even when buyers perceive their output markets as unrelated. Recent work studying information technology markets has argued that changes in industry structure in related markets can have long run implications for product market competition (Bresnahan and Greenstein 1999). However, in contrast to this prior work, we do not require these firms to produce in markets that are substitutes or complements in demand. In our research we argue that changes in structure to markets that share common inputs will have important implications for decisions of quality provisions.

Our paper also makes a contribution to the empirical literature on quality and competition. In general, prior work has demonstrated that increases in competition leads to better quality provision. (e.g. Demberger and Sherr (1989)- legal services industry; Dranove and White (1994) Hospital markets; Borenstein and Netz (1999) and Mazzeo(2003) – airline markets;). While prior work has demonstrated a link between competition and product quality, it has not studied the interaction between technologically related firms as we do.

We also build on existing literature in the area of economics of information security. Losses from attacks are not only influenced by the intensity of attacks, but also on how long the vulnerability remains unpatched (Schneier 2000). As a result, one active area of theory research has examined the economic impacts of vulnerability disclosure and the optimal timing of disclosure (See Arora, Telang and Xu (2004), Cavusaglu et al (2005)). Empirical work examining vulnerability disclosure is rarer. Arora, Nandkumar and Telang (2004) provide empirical evidence on the impact of vulnerability publication when patches are not accompanied by disclosure. Arora, Krishnan, Telang and Yang (2005) use a dataset assembled from CERT/CC and SecurityFocus to show that early disclosure leads to faster patching times. Telang and Wattal (2004) use an event study methodology to show that vulnerability disclosure leads to a loss of market value. However, to our knowledge, no prior work has studied how competition influences vendors' strategic response to vulnerability disclosure.

2. Model

We develop a simple model of firm investment in the time to release of patches for software vulnerabilities. We use this model to develop empirically relevant hypotheses. This model builds upon prior work by Arora, Telang, and Xu (2004). However, in contrast to the prior work, we examine how patching behavior is influenced by market competition, size, and disclosure. The source of competition is a vulnerability that affects multiple products (for future reference we label these *common vulnerabilities*³). The scenario considered by our model is as follows: first a common vulnerability is discovered by an identifier, who informs the intermediary. The intermediary then informs all vulnerable vendors and keeps the vulnerability secret until all the vulnerable vendors fix the vulnerability. The vendor and other vulnerable vendor(s) then commit to a one-time decision on when they will release a patch for the vulnerability. We assume that patches are homogenous and completely remove vulnerabilities. Thus, vendors choose only when to patch and not the quality of the patch that they release.

We allow for stochastic patching times. If x_i is the random variable that denotes the actual patch release time for a vendor i , vendors choose optimal expected patch release

³ A common vulnerability is typically an artifact of a shared code base or design specification or due to a proprietary extension of a widely used software component by a vendor.

time $E[x_i] = \tau_i$, given the distribution of others vendors' patching times. We look for a Nash Equilibrium in which vendors simultaneously pre-commit to an expected patch release date τ , given their expectation of their rivals' τ . Vendors cost function consists of – end user loss θ , and patch development cost, C . End user loss is a function of loss per customer, L , internalization factor, λ and the number of product installations, or quantity q . Thus vendor's cost function is given by

$$V_i = C(\tau_i) + q_i \theta_i (\lambda_i + \omega_i) \quad (1)$$

Let z denote the amount of time end users are exposed to the vulnerability. If there are $J = \{1, 2, 3, \dots, j\}$ other vendors also affected by the vulnerability besides vendor i , then $z \equiv \min\{\tau_1, \tau_2, \dots, \tau_j, s\}$. Let $G(\cdot)$ be the distribution of x_i and let $\Phi(\cdot)$ be the distribution of z . For any vendor i , the expected cost to fix the vulnerability and the expected end user loss are given by $C_i \equiv \int_0^R C(x_i) dG(x_i : \tau_i)$ and $\theta_i \equiv \int_0^R \int_0^{x_i} (L(x_i - z) d\Phi(z : \tau_j, s)) dG(x_i : \tau_i)$ respectively. We assume that C is decreasing and convex in τ , and θ is increasing and convex in τ . It can be shown that the strategies of vendors are complimentary (all proofs are available upon request).

Result 1: *An increase in the number of vendors that share a common vulnerability increases the disclosure threat and leads to declines in the optimal expected patching time τ^* .*

An increase in q increases the expected losses to the vendors and hence results in an earlier patch release by vendors ceteris paribus. Higher λ (which happens when vendor j is also a rival) results in a pre-commitment to patch earlier. In this paper, we refer to the effect of λ as the effect of competition.

Result 2: *An increase in the number of rivals that share a common vulnerability increases λ and leads to declines in the optimal expected patching time τ^* .*

Result 3: *τ^* for a vendor is decreasing in quantity.*

3. Data and variables:

We assembled our data set from publications of CERT/CC. CERT is a federally funded organization whose main task is to disseminate vulnerability information. We obtained data from CERT/CC over the period September 2000 to August 2003. We further augmented this data with measures of quantity (collected from Harte Hanks CI database) and competition. Our measure of competition consists of 3 separately constructed measures: *VENDORS* is equal to the total number of vendors listed as “vulnerable” by CERT for a specific vulnerability. *RIVALS* is equal to the number of vendors that CERT list as vulnerable and that operate in the same market as the vendor in the vendor-vulnerability pair. *NONRIVALS* is equals to the number of vendors that are vulnerable but which operate in different markets. In all, the sample consists of 241 distinct vulnerabilities and 461 observations for 16 different “closed source” vendors. We also include a measure to capture the severity of vulnerability (assigned by CERT/CC) as *LOGSEVERITY*.

Our dependent variable is *DURATION*, which measures the elapsed time in calendar days from the date when the vendor knew of the vulnerability until the patch release date. The value of *DURATION* depends on– *instant or non-instant disclosure*. If the vulnerability is *instantly disclosed*, *DURATION* is the elapsed time in days between when the vulnerability was publicly disclosed and the time the vulnerability was patched by the vendor. If the vulnerability was *non-instantly disclosed*, *DURATION* is the elapsed time between when CERT/CC informed the

vendor of the existence of the vulnerability and when the vendor issued a patch. For the empirical analysis we use *LOGDURATION* as our dependent variable.

4. Statistical Method and Identification

Our goal is to examine how the patching times for vendor i in market m facing vulnerability v varies with competition, disclosure, and quantity. If Z and X denotes vulnerability and vendor characteristics respectively, then

$$LOGDURATION_{imv} = \beta_0 + \beta_1 COMPETITION_{imv} + \beta_2 DISCLOSURE_v + \beta_3 LOGQUANTITY_i + \theta_1 X_i + \theta_2 Z_v + \varepsilon_{iv} \quad (2)$$

To empirically separate the effects of competition and disclosure on vendor patching times, we utilize two sources of variance in our data. First, we utilize variance in the number of rivals and nonrivals affected by the vulnerability. Increases in the number of direct rivals to the vendor will influence patching times through both the competition and disclosure effects, while increases in nonrivals will influence patching times only through disclosure. Under this method we use $RIVAL S_{iv}$ to proxy for $COMPETITION_{iv}$ and $RIVAL S_{iv} + NONRIVAL S_{iv}$ to proxy for $DISCLOSURE_{iv}$, giving us the following

$$LOGDURATION_{imv} = \beta_0 + \beta_1 RIVAL S_{imv} + \beta_2 (RIVAL S_{imv} + NONRIVAL S_{imv}) + \beta_3 LOGQUANTITY_i + \theta_1 X_i + \theta_2 Z_v + \varepsilon_{iv} \quad (3)$$

We then utilize variance in how vendors are informed of vulnerabilities. Vulnerabilities are *publicly disclosed* when a third party or another vendor announces the existence of a vulnerability, and they are *privately disclosed* when CERT/CC informs the vendor of the presence of a vulnerability while the vulnerability remains unknown to the general public. We identify the *competition* effect by examining how changes in the number of affected vendors influence patching time when vulnerabilities are publicly disclosed. We identify the *disclosure* effect by comparing how changes in the number of affected vendors influence average patching times under private and public disclosure. Thus, our second model utilizes variation in vulnerability disclosure to place bounds on the structural parameters. The effect of marginal vendor on *disclosure* is zero under instant disclosure since the vulnerability has already been disclosed. Thus, we can decompose the effect of number of other vendors on patching times as follows:

$$LOGDURATION_{imv} = \beta_0 + \beta_1 VENDORS_{imv} + \beta_2 (1 - INSTANT_{imv}) * VENDORS_{imv} + \beta_3 LOGQUANTITY_i + \lambda INSTANT_{imv} + \theta_1 X_i + \theta_2 Z_v + \mu_v + \varepsilon_{iv}$$

Under this method we use $VENDORS_{imv}$ to proxy for $COMPETITION_{iv}$ and $(1 - INSTANT_{imv}) * VENDORS_{imv}$ to proxy for both $DISCLOSURE_{iv} + COMPETITION_{iv}$. We use differences in these estimates to place bounds on the effect of disclosure threat ($\beta_2 - \beta_1$). Our results summarized below imply that higher amounts of competition, disclosure threat and quantity result in earlier patch releases on an average.

Table 8 Summary of Results: Elasticity Estimates Across Models for 10% increase in competition, disclosure & quantity

	Identification using Rivals and Nonrivals (2)	Identification using Instant Disclosure (3)
Competition Effect	-2.7%	-5.4%
Disclosure Effect	-2.1%	-0.9%

Disclosure Effect + Competition Effect	-4.8%	-6.3%
Quantity Effect	-1.3%	-1.4%

Stated in number of days, a 10% increase in the number of competitors lowers expected patching times between 7 to 25 days due to competition. Further a 10% increase in the number of such vendors lowers expected patching times between 12 to 15 days due to disclosure threat. Also, a 10% increase in installed base is associated with an earlier patch release by about 2.5 days.

5. Discussion and conclusion

This research provides evidence on how competition influences quality provision in information technology markets. We provide a framework for understanding how vendors in one market may influence quality provision in another. Further, in contrast to prior research, we show that such linkages can be important even when vendors operate in unrelated output markets. These results also have implications for the debate of how to improve software quality. Our research demonstrates that despite high levels of concentration in many software markets, threat of disclosure from vendors in complementary markets works to reduce patching times almost as much as increases in the number of competitors. Our results also suggest that *non-instant* disclosure could be more welfare-enhancing than *instant disclosure*.

References:

- Arora A., Krishnan R, Telang R. & Yang Y. (2005) "An Empirical Analysis of Vendor Response to Disclosure Policy," 4th WEIS, Harvard University, Boston MA.
- Arora A., Nandkumar A. & Telang R. (2004) "Impact of patches and software vulnerability information on frequency of security attacks - An empirical analysis," Working paper, CMU.
- Arora A., Telang R. & Xu H. (2004) "Optimal Policy for Software Vulnerability Disclosure," 3rd WEIS University of Minnesota, MN.
- Borenstein S. and Netz J. (1999), "Why do All the Flights Leave at 8 am?: Competition and Departure-Time Differentiation in airline markets," *International Journal of Industrial Organization*, 20(3):344-365.
- Bresnahan, T., and S. Greenstein, (1999), "Technological Competition and the Structure of the Computer Industry," *Journal of Industrial Economics*, 47(1): 1-40
- Dranove D. and W.White (1994), "Recent Theory and Evidence on Competition in Hospital Markets," *Journal of Economics and Management Strategy*, 3(1):169-209.
- Domberger S. and A. Sherr (1989), "The impact of competition on pricing and Quality of Legal Services," *International Review of Law and Economics*, 9:41-56.
- Mazzeo M. (2003), "Competition and Service Quality in the U.S. Airline Industry," *Review of industrial Organization*, 22: 275-296
- Schneier B. (2000) "Full Disclosure and the Window of Exposure," in: *CRYPTO-GRAM*, 2000.
- Telang R. and Wattal S. (2005) "Impact of Software Vulnerability Announcements on the Market Value of Software Vendors – an Empirical Investigation," 4th WEIS, Harvard University, Boston, MA,