An Economic Analysis of Market for Software Vulnerabilities

Karthik Kannan* Rahul Telang †

May 3, 2004

Abstract

Software vulnerability disclosure has become a critical area of concern for policy-makers. Traditionally, Computer Emergency Response Team (CERT) has been acting as an infomediary between benign identifiers (who report vulnerability information voluntarily) and software users. After verifying a reported vulnerability, the infomediary – CERT – sends out a public "advisory" so that users can safeguard their systems against potential exploits. Of late, firms such as iDefense have been implementing a different market-based approach for vulnerability disclosure where the "market-based" infomediary provides monetary rewards to identifiers for each vulnerability disclosed to it. The infomediary shares this information with its client base. Using this information, clients protect themselves against attacks that exploit those specific vulnerabilities.

The key question addressed in our paper is whether movement towards such a market-based mechanism for vulnerability disclosure leads to a better social outcome. Our analysis demonstrates that an active "market-based mechanism" for vulnerabilities almost always underperforms a passive CERT-type mechanism. We provide intuitions to this counter-intuitive result. Further, our paper provides policy recommendations that improve the relative performance of the market-based mechanism though not completely. Finally, we extend our analysis and analyze a new mechanism – "Federally-Funded Social Planner" – that always performs better than a market-based mechanism.

1 Introduction

One of government's fundamental jobs is deciding what goods and services should be provided by which types of markets. The US has decided that postal delivery and national defense services should be provided by the government. Utilities used to be primarily regulated monopolies but now operate in regulated competition. Grocery stores are largely unregulated. The choice is usually (supposed to be) made on the basis of social welfare, including efficiency and equity considerations. Here we offer the first ever such analysis with regard to the market for software vulnerability detection.

Traditionally, Computer Emergency Response Team (CERT) has been acting as an *infomediary* between *benign identifiers* who report vulnerability information and software users. CERT's role evolved during the early days of the Internet when vulnerability discovery and reporting was relatively infrequent.

*Purdue University, kkarthik@mgmt.purdue.edu

†Carnegie Mellon University, rtelang@andrew.cmu.edu

Since no market existed for vulnerabilities, CERT's role was crucial in disseminating vulnerability information efficiently. After verifying a reported vulnerability and coordinating with vendors, the infomediary – CERT – typically sends out a public "advisory" so as to allow users to safeguard their systems against potential exploits. In order to ensure that such public notifications are not exploited by *attackers* to attack software users, CERT follows a series of steps before such a disclosure. The steps include contacting the vendor for the appropriate patch, waiting for an appropriate time before publicly disclosing the vulnerability, etc. In this mechanism, reporting vulnerabilities is voluntary with no explicit monetary gains to benign identifiers.

Of late, either due to the increased value of electronic assets or due to the penetration of the Internet, the number of vulnerabilities discovered has increased. For example, 4129 vulnerabilities were reported in 2002, whereas only 1090 were reported in 2000 (CERT (2003)). This has also led to the creation of a market for vulnerability where firms such as iDefense[1] have been acting as infomediaries. In this *market-based mechanism*, the infomediary offers monetary reward to the identifier for every vulnerability reported to it. The infomediary then shares this information with users of this software who are subscribed to its service. Using this information along with other value added services[2] which the market-based infomediary provides, subscribers can protect themselves against attacks that exploit those specific vulnerabilities.

The key question addressed in this paper is whether such a movement towards a market-based mechanism leads to a better social outcome? The answer is not obvious. On one hand, monetary incentives to discover vulnerabilities may lead to benign identifiers investing more effort and time in finding them. And this would lead to a better social outcome. But on the other hand, the same incentives may also lead to a *race* for vulnerability discovery (See Dasgupta & Stiglitz (1980) for race in R&D investments) between benign identifiers and attackers. If this happens and the number of vulnerabilities discovered by the attacker increases, it may decrease social welfare. Note that a monopolistic market-based infomediary has an incentive to serve only a fraction of the entire market exposing users not subscribed to infomediary's service to attacks. Also, the non-subscribers may suffer adversely when the monopolist misuses vulnerability information to increase its profits. This may lead to further destroying social welfare. We term such a market as an *unregulated market*. But even in a *regulated market* where a monopolistic infomediary cannot misuse information, the answer is unclear.

From a policy-maker's perspective, understanding this question is crucial. If markets perform at least as well as traditional mechanisms like the CERT-type one, then policy-makers need to reshape the role of such institutions in the future. Moreover, this also means that our policies should encourage such markets. On the other hand, if markets decrease welfare and they are here to stay, then policy-makers need to think about regulations that may achieve the desired objective. One key contribution of our paper is to argue that while software security typically has been a domain of computer scientists and technical researchers, it is the emerging economic and market mechanism questions that have significant social welfare implications. Even so, there is little academic research in this area to draw from. Our paper tries to bridge this gap by analyzing the economic efficacy of these mechanisms and by providing appropriate policy guidelines.

[1] www.iDefense.com/contributor.html

[2] For example, the infomediary may deliver a patch for the vulnerability or provide filters to protect against attacks that exploit the vulnerability. Sometimes, simply the information that such vulnerability exists is crucial for firms because they usually can not keep track of all such details.

2 Literature Review

Much of the prior work in the software vulnerability and information security area has focused on the technical aspects of the problem (e.g., Krsul *et al.* (1998), Du & Mathur (1998)). There have software engineering literature that has focused on process improvements and shown that when such techniques are incorporated, the software quality improves i.e., the software has fewer vulnerabilities (Banker *et al.* (1998)). But it is widely believed that vulnerabilities and therefore, attacks exploiting these vulnerabilities cannot be completely eliminated.

Given this, a few papers have analyzed related problems in the information security space (in the interest of space, we list those relevant papers). Gordon & Loeb (2002) develop an economic model for information security investment decisions. Similarly, Arora *et al.* (2003) develop an economic model to study a vendor's decision of when to introduce its software and whether or not to patch vulnerabilities in its software. To our knowledge, no prior work has addressed specific issues discussed in the introduction. Academicians and practitioners in different capacities have been proposing different legal/economic frameworks for software vulnerability disclosure (e Week (2003), Varian (2000)). But since this area of research is relatively nascent and much of the work is yet to come, policy-makers are left with little guidance in understanding the implications of different frameworks. Our paper is unique in providing a formal model to analyze different disclosure mechanisms.

3 Model

There are four main participants in our model – the infomediary,[3] a benign identifier, a attacker and software users. Note that we consider only a monopolistic market-structure for the infomediary. It is interesting because this market is likely to yield to natural monopoly. In this market, let p_b be the reward paid to the benign identifier by the infomediary for each vulnerability reported. Let p_s represent the one-time subscription fee that the infomediary charges each of its subscribers – software user.

The (p_b, p_s) pair set by the infomediary determines the number of subscribers (and hence the market share) to this service, the number of vulnerabilities reported by the benign identifier and the probability of attacks. But the optimal prices p_b and p_s are themselves determined by market share, number of vulnerabilities reported, etc. Therefore, we model this as a two period game. In the first period, the infomediary sets its pricing policy to maximize its objective function and in the second period, all other players – software users, the benign identifier and the attacker – react. However when solving this game, we first solve for the reaction of the benign identifier, the attacker and software users for a given (p_b, p_s) pair and then, solve for the optimal (p_b, p_s) using backward induction. Ultimately, our goal is to calculate the welfare-metrics – the total industry loss and the total user loss – for each mechanism based on these prices and market share.

Without loss of generality, we assume that there is one vulnerability in the product and that the benign identifier and the attacker attempt to discover it. Having only one vulnerability allows us to model everything as probability measures. Let K_{attacker} be the probability that the vulnerability is first discovered by the attacker. In this case, the attacker exploits the vulnerability to attack all users (even including the subscribers to the infomediary's service). Similarly, let K_{reported} be the probability that the vulnerability is first discovered by

[3] iDefense, CERT, federally funded planner are all examples of infomediary.

the benign identifier who reports it to the infomediary[4]. After obtaining the vulnerability information, the infomediary notifies its subscribers so that they can protect their systems against potential future attacks. Let $K_{\text{prevented}}$ represent the probability that the attack is prevented by subscribing to the infomediary's service i.e., $K_{\text{prevented}}$ is the value provided by the infomediary. This probability is distinguished between a regulated mechanism and an unregulated mechanism.

The key consideration here is what does the infomediary do with the vulnerability information? We argue that the market-based infomediary always has an incentive to "leak" vulnerability information without proper safeguard. This serves to threaten to non-subscribers who may be subjected to attacks exploiting the vulnerability. Such a mechanism is referred to as the unregulated mechanism. In contrast, the monopolistic infomediary in a regulated market will make the information public only with proper safeguards such that users not subscribed to the service are not affected adversely. Therefore, we use superscripts to distinguish between the regulated and the unregulated case.

$$K_{\text{prevented}} = \begin{cases} K^{\text{leak}}_{\text{prevented}} & \text{if it is an unregulated market} \\ K^{\text{no leak}}_{\text{prevented}} & \text{if it is a regulated market} \end{cases}$$

4 Unregulated Market

Let N represent market share of infomediary. Our objective, in this subsection, is to characterize the expressions for the probabilities $K_{\text{reported}}, K^{\text{leak}}_{\text{prevented}}, K_{\text{attacker}}$ and N as functions of p_b and p_s

4.1 Characterizing the Number of Subscribers

We assume that software users are heterogeneous in terms of the loss they incur when a vulnerability is exploited. Let the user "loss"-type, θ, be distributed between $[0, \overline{\theta}]$ with the distribution function $F(\theta)$. Any software user i of type θ_i is assumed to incur a loss of θ_i^2 when the vulnerability is exploited. The software users have an option of preventing attacks on their systems by subscribing to the infomediary's service. Let the subscription fee charged by infomediary be p_s. Any user i, whose expected profit from subscribing, $\Pi_{\text{user}} = \theta_i^2 \; K^{\text{leak}}_{\text{prevented}} - p_s > 0$ subscribes to the service. The first term corresponds to the loss prevented by subscribing to the service and the second term corresponds to the payment made to the infomediary. Given these, infomediary's market share is:

$$N = 1 - F\left(\sqrt{\frac{p_s}{K^{\text{leak}}_{\text{prevented}}}}\right) \tag{1}$$

In a mechanism where software users are not charged any price at all i.e., $p_s = 0$, $N = 1$.

4.2 Characterizing K_{reported}, $K^{\text{leak}}_{\text{prevented}}$, and K_{attacker}

The (p_b, p_s) pair set by the infomediary determines the effort levels exerted by the benign identifier and the attacker which in turn, dictate K_{reported}, $K^{\text{leak}}_{\text{prevented}}$ and K_{attacker}. Although our results hold good for any generic functional forms, we need specific expressions to be able to obtain a tractable solution. We obtain this

[4]By definition, a benign identifier does not exploit the vulnerability.

expression by modeling the competition between a benign identifier and a attacker within the software's life cycle period, T.

Competition between the Benign Identifier and the attacker

Within the time period T, we assume that the probability that a *player* – the benign identifier or the attacker – discovers the vulnerability without exerting any effort is a random variable $\gamma \in [0, 1]$ with uniform distribution. Thus, given our distributional assumption, $\frac{\gamma}{T}$ is the probability density function (pdf) for the vulnerability being discovered by either player at any time $t < T$. Players can alter γ and hence, the pdf by exerting effort. We assume that a benign identifier exerts an effort α. This effort increases its pdf to $\frac{\alpha+\gamma}{T}$. Similarly, the attacker exerts an effort level of β that increases its pdf to $\frac{\beta+\gamma}{T}$. Given these effort levels, we compute the probabilities"
$\underline{K_{\text{reported}}}$: The probability that the vulnerability is reported corresponds to the probability that the vulnerability is first discovered by the benign identifier and reported to the infomediary.

$$K_{\text{reported}} = \int_0^T \text{Probability}(benign = t)\ \text{Probability}(\overline{attacker} < t)\ \text{d}t$$

Probability$(benign = t)$ is the probability that the vulnerability is identified by the benign identifier at time t by exerting an effort α. Probability$(\overline{attacker} < t)$ is the probability that the vulnerability has *not* been identified by the attacker exerting effort β until time t. Therefore,

$$K_{\text{reported}} = \int_0^T \frac{\alpha+\gamma}{T}\left(1 - \frac{(\beta+\gamma)t}{T}\right)\text{d}t = (\alpha+\gamma)\left(1 - \frac{(\beta+\gamma)}{2}\right) \quad (2)$$

$\underline{K_{\text{prevented}}}$:The probability that an attack is prevented corresponds to the value provided by the infomediary's service. It is important to note that when the infomediary in an unregulated market "leaks" the vulnerability information without proper safeguards, all reported vulnerabilities become exploitable. Thus by subscribing to infomediary's service, a user can prevent all those attacks that occur whenever the benign identifier reports the vulnerability to the infomediary. Therefore, the value of infomediary services is simply

$$K^{\text{leak}}_{\text{prevented}} = K_{\text{reported}} = (\alpha+\gamma)\left(1 - \frac{(\beta+\gamma)}{2}\right) \quad (3)$$

We will show in the next section how to calculate $K^{\text{no leak}}_{\text{prevented}}$ when the monopolist is regulated.
$\underline{K_{\text{attacker}}}$: the probability that the vulnerability is first discovered by the attacker is

$$\begin{aligned} K_{\text{attacker}} &= \int_0^T \text{Probability}(attacker = t)\ \text{Probability}(\overline{benign} < t)\ \text{d}t \\ &= (\beta+\gamma)\left(1 - \frac{(\alpha+\gamma)}{2}\right) \end{aligned} \quad (4)$$

When the vulnerability is first discovered by the attacker, the attacker attacks all users including the subscribers of the infomediary's service.

Optimal Effort Level

Recall that the effort exerted by the benign identifier increases her probability of finding the vulnerability to $\alpha + \gamma$. This effort is rewarded with p_b if she discovers the vulnerability before the attacker. Since K_{reported} is the probability that the benign identifier discovers the vulnerability first, her expected revenue is $p_b\ K_{\text{reported}}$. For some effort level α, the benign identifier's cost is $C(\alpha)$. Thus, the expected profit for the benign identifier is $\Pi_b = K_{\text{reported}}\ p_b - C(\alpha)$. If $C(\alpha) = M\alpha^2$ where M is an exogenous constant used for scaling purposes such that α, and β are bounded $[0, 1]$,

$$\Pi_b = (\alpha + \gamma)\left(1 - \frac{(\beta + \gamma)}{2}\right) p_b - M\ \alpha^2 \tag{5}$$

Next, consider the attacker's expected profit. The attacker benefits by attacking all users if he discovers the vulnerability first. But if he discovers the vulnerability after the benign identifier, he obtains the profit only from attacking users not part of the infomediary's subscription.[5] We assume that if the attacker is successful in attacking a user of type θ_i, he gains a profit of θ_i. Note that the functional form of the attacker's profit function is intentionally made to be different from the loss suffered by the user – θ_i^2.[6] The attacker's cost is $C(\beta)$. Therefore,

$$\Pi_h \quad = \quad \underbrace{K_{\text{attacker}}\left(\int_0^{\overline{\theta}} \theta\ \mathrm{d}F(\theta)\right)}_{\text{Attacker discovers first: Attacks all users}} + \underbrace{K^{\text{leak}}_{\text{prevented}}\left(\int_0^{\sqrt{\frac{p_s}{K^{\text{leak}}_{\text{prevented}}}}} \theta\ \mathrm{d}F(\theta)\right)}_{\text{Attacker discovers next: Attacks non-subscribers}} - C(\beta)$$

The optimal attacker effort, β^*, is a solution of this implicit equation which requires some functional form assumption for $F(\theta)$. To ensure analytical tractability, we let θ to be distributed uniformly $[0, \overline{\theta}]$. It means, $F(\theta) = \frac{\theta}{\overline{\theta}}$. Note that this assumption, when combined with the non-linear loss function – θ_i^2 – assumed for each user, reflects the empirical observations quite well i.e., many users suffer smaller losses while a few users suffer huge losses. Substituting for $F(\theta)$ and simplifying the equation,

$$\Pi_h \quad = \quad (\beta + \gamma)\left(1 - \frac{(\alpha + \gamma)}{2}\right)\frac{\overline{\theta}}{2} + K^{\text{leak}}_{\text{prevented}}\frac{p_s}{K^{\text{leak}}_{\text{prevented}}\ 2\overline{\theta}} - M\ \beta^2$$

Since we know the expected profits for both benign identifier and the attacker, we can solve for optimal α and β and get

$$\begin{aligned} \alpha^* &= \frac{(8\ M - \overline{\theta})\ p_b\ (2 - \gamma)}{32\ M^2 - p_b\ \overline{\theta}} \\ \beta^* &= \frac{(2 - \gamma)(4\ M - p_b)\overline{\theta}}{32\ M^2 - p_b\ \overline{\theta}} \end{aligned}$$

[5] It is trivial to show that attacker never finds it optimal to sell the vulnerability.

[6] In some cases, the attackers may gain a lot by gaining access even though users may not lose a lot. In some other cases, the attackers may not gain as much but the cost to the user could be really high. For example, sometimes, the attackers may take down a web-site causing significant damages to users though they might not gain correspondingly.

Note that since $\alpha + \gamma$ and $\beta + \gamma$ are probabilities, they should be bounded $[0, 1]$ for any reasonable result. We let the scaling factor M to ensure this. Let M_{th} be the threshold value above which the probabilities are bounded. For the rest of the analysis, we assume $M > M_{th}$. For $M > M_{th}$, we observe the following properties in these equations:

- Both parameters – α and β – are independent of p_s.
- As p_b increases, α increases but β decreases. This suggests that while effort exerted by the benign identifier increases with p_b, this, in turn, imposes a negative externality on the attacker's incentives and reduces his efforts.
- For a given p_b, both the benign identifier and the attacker have incentives to increase their efforts as γ decreases.
- Finally, as M increases, i.e., the cost of exerting effort increases, the optimal effort levels – α^* and β^* – decrease as expected.

Based on the values, one can now calculate various probabilities. We note that $\frac{\partial K_{\text{reported}}}{\partial p_b} > 0$; and $\frac{\partial K_{\text{attacker}}}{\partial p_b} < 0$. Therefore, attackers and benign identifiers impose negative externality on each other. Moreover, as the baseline probability of discovering the vulnerability without effort – γ – increases, all three probabilities increase i.e., $\frac{\partial K_{\text{reported}}}{\partial \gamma} > 0$, $\frac{\partial K^{\text{leak}}_{\text{prevented}}}{\partial \gamma} > 0$, $\frac{\partial K_{\text{attacker}}}{\partial \gamma} > 0$. Finally M increases, all three probabilities decrease i.e., $\frac{\partial K_{\text{reported}}}{\partial M} < 0$, $\frac{\partial K^{\text{leak}}_{\text{prevented}}}{\partial M} < 0$, $\frac{\partial K_{\text{attacker}}}{\partial M} < 0$.

4.3 Optimal Pricing p_b and p_s

As is common in subgame perfect equilibrium, we first calculate the second period consequence of first period action and based on those outcomes, calculate the optimal first period actions. Now based on the calculated probabilities, we calculate the optimal p_s and p_b which the market-based infomediary sets. The infomediary maximizes

$$\max_{p_b, p_s} \quad \underbrace{N\, p_s}_{\text{Revenue to the infomediary}} \quad - \quad \underbrace{K_{\text{reported}} p_b}_{\text{Cost incurred to pay for each vulnerability reported}}$$

We substituting for N from equation 1 and using $F(\theta) = \frac{\theta}{\overline{\theta}}$. Based on the first order conditions, we derive p_s^* and p_b^*. We also observe that $\frac{\partial p_b^*}{\partial \gamma} < 0$ which implies that as users voluntarily provide vulnerability information, incentives to find vulnerability disclosure decreases.

4.3.1 Welfare-Metrics

Our final goal is to analyze how the welfare changes under different market conditions. To measure the efficacy of this unregulated market-based mechanism, we compute the total user loss as:

$$UL^{\text{leak}}_{\text{MARKET}} = \underbrace{K_{\text{attacker}} \left(\int_0^{\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right)}_{\text{Attacker discovers first: Attacks all users}} + \underbrace{K^{\text{leak}}_{\text{prevented}} \left(\int_0^{(1-N)\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right)}_{\text{Attacker discovers next: Attacks non-subscribers}} + \underbrace{N\, p_s}_{\text{Subscription Paid}} \tag{6}$$

By substituting for p_s^*, p_b^*, $K_{\text{prevented}}^{\text{leak}}$ and K_{attacker}, one can compute $UL_{\text{MARKET}}^{\text{leak}}$. Similarly, one can compute the total industry loss. Equation 6, which corresponds to the user loss, is combined with the infomediary's profit to obtain the industry loss expression

$$IL_{\text{MARKET}}^{\text{leak}} = K_{\text{attacker}} \left(\int_0^{\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + K_{\text{prevented}}^{\text{leak}} \left(\int_0^{(1-N)\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + K_{\text{reported}} \, p_b \tag{7}$$

Given these expressions, the following observation is worth noting: For a given $\overline{\theta}$, p_b, p_s and M, recall that $K_{\text{prevented}}^{\text{leak}}$, K_{attacker} and K_{reported} increase as γ increases. But as γ increases, p_b^* decreases which further aids the increase in K_{attacker} (since $\frac{\partial K_{\text{attacker}}}{\partial p_b} < 0$). Both these factors make $UL_{\text{MARKET}}^{\text{leak}}$ and $IL_{\text{MARKET}}^{\text{leak}}$ increase with γ.

4.4 User Loss in CERT-Type Mechanism

Recall that in the CERT-type mechanism, no money is paid to the benign identifier for reporting the vulnerability, i.e., $p_b = 0$. Also, no subscription is charged and the vulnerability information is provided to all users i.e., $p_s = 0$ and $N = 1$. Given this, the user loss and the industry loss are identical in the CERT-type mechanism:

$$UL_{\text{CERT}} = IL_{\text{CERT}} = K_{\text{attacker}} \left(\int_0^{\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) \tag{8}$$

To compute this equation, we derive the expression for α^* and β^* using a framework similar to that in the earlier section. $\alpha^* = 0$ and the benign identifier does not exert any effort at all. But the vulnerability is still discovered by the benign identifier with a probability of γ, and reported to the infomediary. On the other hand, the attacker invests an optimal β^*. Using α^* and β^*, we compute K_{attacker} which is then substituted back in equation 8 to obtain

$$UL_{\text{CERT}} = IL_{\text{CERT}} = \frac{(2-\gamma)((2-\gamma)\overline{\theta} + 8\,M\gamma)\overline{\theta}^2}{48\,M} \tag{9}$$

When $\gamma = 0$, $UL_{\text{CERT}} = IL_{\text{CERT}} = \frac{\overline{\theta}^3}{12\,M}$. This corresponds to the condition when the vulnerability is never reported to the CERT-type infomediary. But as γ increases, the CERT-type infomediary provides some value. This is because as γ increases, the probability that the benign identifier reports the vulnerability is higher. But the same is true for the attacker. The attacker also finds it easier to discover the vulnerability which implies that the probability of an attack exploiting the vulnerability increases. Hence the higher the γ, the higher the user loss. But the rate of increase of UL_{CERT} and IL_{CERT} with respect to γ is lower in the CERT-type mechanism than in the market-based mechanism. This drives the comparisons executed in the following subsection.

4.5 Comparative Static: CERT Versus Unregulated Market

The following propositions outline the main insight (Proofs have not been shown because of space constraints):

Proposition 4.1 *1. Even at $\gamma = 0$, for a given M, there exists a $\overline{\theta}$ such that the user loss in the unregulated market-based mechanism is more than that in the CERT-type one.*

2. *At $\gamma = 0$, for some $M > \hat{M}$, the user loss in the unregulated market-based mechanism is always more than that in the CERT-type mechanism.*

The striking part of the result is that even when $\gamma = 0$, the market-based mechanism may underperform relative to its CERT-type counterpart. Note that since no one reports any vulnerability information voluntarily to CERT when $\gamma = 0$, CERT has no role to play. In short, there is no market left. But even when $\gamma = 0$, the market-based infomediary gathers vulnerability information from the benign identifier by rewarding discovery and disseminates that information to its subscribers. In other words, an active market exists. One would expect that having even a monopolistic market-based infomediary is better than having none at all. But our results show that a monopolistic market-based infomediary in an unregulated market is almost always worse than having no market at all from users' point of view.

What is the intuition behind this perverse result? The key insight is that a monopolistic market-based infomediary in an unregulated framework always has an incentive to misuse the vulnerability information. Whenever the benign identifier reports the vulnerability information, the infomediary protects its own subscribers but it "leaks" the information without appropriate safeguards. This "leakage" exposes non-subscribers to attacks from the attacker. The "leakage" also serves to increase the users' incentives to subscribe to the infomediary's service which, in turn, allows the monopolist to charge a higher subscription fee, p_s, and thus eroding user welfare.

These two propositions highlight the fact that an unregulated market-based mechanism will be better than the CERT-type mechanism only for a small parameter region. Otherwise, the unregulated market-based mechanism is worse than a "no market" mechanism like the CERT-type one. Stated differently, doing nothing to incentivize vulnerability discovery is almost always better than letting a monopolist enter an unregulated market.

5 Regulated Market – Without "Leakage"

As we noted previously, the key difference between a regulated and an unregulated market-based mechanism is whether the infomediary "leaks" the vulnerability information. In an unregulated market-based mechanism, the infomediary "leaks" the vulnerability information without proper safeguard. Hence, subscribing to the infomediary's service prevents the attack that would have resulted even when the benign identifier reports the vulnerability to infomediary. Therefore, $K^{\text{leak}}_{\text{prevented}} = K_{\text{reported}}$ as shown in equation 3. But in the regulated market-based mechanism, the infomediary is bounded by regulation to not "leak" the vulnerability information without safeguards. In this case, the value of the infomediary's service, $K^{\text{no leak}}_{\text{prevented}}$, is simply the number of vulnerabilities discovered by the benign identifier that could have otherwise resulted in attacks. Mathematically,

$$K^{\text{no leak}}_{\text{prevented}} = \int_0^T \text{Probability}(attacker = t)\ \text{Probability}(benign < t)\ \text{d}t = (\alpha + \gamma)\left(\frac{\beta + \gamma}{2}\right)$$

The other probabilities K_{attacker} and K_{reported} remain the same. Similar to the unregulated mechanism, we derive the expressions for α^* and β^* which are then substituted to compute K_{attacker}, K_{reported}, and $K^{\text{no leak}}_{\text{prevented}}$. Similar to the earlier case, the infomediary maximizes its expected profit function

$$\max_{p_b, p_s} N\ p_s - K_{\text{reported}} p_b \tag{10}$$

Based on the first order conditions, we obtain p_s^* and p_b^*. Using these p_b^*, p_s^*, $K_{\text{prevented}}^{\text{no leak}}$ and K_{attacker}, we also calculate the total user loss, $UL_{\text{MARKET}}^{\text{no leak}}$, and the total industry loss, $IL_{\text{MARKET}}^{\text{no leak}}$.

$$UL_{\text{MARKET}}^{\text{no leak}} = K_{\text{attacker}} \left(\int_0^{\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + K_{\text{prevented}}^{\text{no leak}} \left(\int_0^{(1-N)\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + N \, p_s \tag{11}$$

$$IL_{\text{MARKET}}^{\text{no leak}} = K_{\text{attacker}} \left(\int_0^{\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + K_{\text{prevented}}^{\text{no leak}} \left(\int_0^{(1-N)\overline{\theta}} \frac{\theta^2}{\overline{\theta}} \, \mathrm{d}\theta \right) + K_{\text{reported}} \, p_b \tag{12}$$

Note that the expressions are similar to those in equation 6 and equation 7 except that we use $K_{\text{prevented}}^{\text{no leak}}$ instead of $K_{\text{prevented}}^{\text{leak}}$. We are again interested in comparing the performance of the regulated market-based mechanism with the CERT-type mechanism. Similar to the earlier case, note that the rate of increase of user loss with respect to γ is higher in the regulated market-based mechanism than in the CERT-type one.

5.1 Comparative Static: CERT versus Regulated Market

The following proposition shows that the performance of the market-based scheme improves but only marginally.

Proposition 5.1 *There always exists a $\gamma' > 0$ below which a regulated market-based mechanism outperforms the CERT-type one. Otherwise, the CERT-type mechanism is better.*

Reassuringly, at least we find that when $\gamma = 0$, a regulated market-based mechanism outperforms the CERT-type mechanism. We reiterate that when $\gamma = 0$, no vulnerabilities are reported to the CERT-type infomediary and therefore, the CERT-type mechanism has little value. In contrast, the market-based mechanism incentivizes the benign identifier to discover the vulnerability. Since the regulation prevents the monopolistic market-based infomediary from misusing the information, we observe that the market-based scheme outperforms the CERT-type one (which is a "no market" mechanism). Therefore, the idea that even a monopolist is better than having no market at all, holds true in this case.

As γ increases, both the CERT-type mechanism and the market-based mechanism incur higher loss. However, recall that the rate of increase of user loss in the market-based mechanism is higher than that in the CERT-type mechanism. This implies that markets are better only for some lower value of γ. Beyond the critical value of γ, even the regulated market-based mechanism underperforms.

6 Is there a better Mechanism? Federally Funded Social Planner

The major goal of this paper is to analyze the welfare implications of different software vulnerability disclosure mechanisms. In this section, we extend the earlier model to investigate whether a better mechanism exists. We determine that the optimal mechanism is the one where an infomediary incentivizes vulnerability discovery but it does not charge any subscription fee and provides vulnerability information to all users with proper safeguards. Stated differently, the optimal mechanism is akin to a "Federally Funded" program where an infomediary like CERT acts not only as a social planner in minimizing the industry loss but also as an infomediary incentivizing vulnerability discovery by paying p_b.

Recall that the total industry loss is given by

$$IL = \underbrace{K_{\text{attacker}}\left(\int_0^{\overline{\theta}} \theta^2 d\theta\right)}_{\text{Attacker discovers first: Attacks all users}} + \underbrace{K^{\text{leak}}_{\text{prevented}}\left(\int_0^{\sqrt{\frac{p_s}{K_{\text{prevented}}}}} \theta^2 d\theta\right)}_{\text{Attacker discovers next: Attacks non-subscribers}} + \underbrace{K_{\text{reported}}\, p_b}_{\text{Money to benign identifiers}}$$

The infomediary's objective function is

$$\min_{p_b, p_s} IL \tag{13}$$

In order to solve for the optimal p_b and p_s, we obtain the expression for the probabilities – K_{attacker} and K_{reported} – in a manner similar to that of the unregulated market. Based on the first order condition, we obtain $p_s^* = 0$. This means that disclosing the vulnerability to all users is the optimal mechanism. Similarly, we obtain the optimal p_b^* as well.

As expected, we note that p_b^* decreases as γ increases. This is intuitive – as the benign identifier finds it easier to search and report vulnerabilities, there is little reason to encourage. From this, it is immediately obvious that for some value of γ, $p_b = 0$ and that corresponds to the federally funded mechanism being identical to the CERT-type mechanism. We compute the threshold value of γ when the CERT-type mechanism is identical to the Federally funded mechanism as

$$\overline{\gamma}_{\text{FED}} = \frac{2\overline{\theta}^3}{48\ M^2 - 4M\ \overline{\theta}^2 + \overline{\theta}^3}$$

Naturally, our interest lies in comparing the two schemes when $\gamma \leq \overline{\gamma}_{\text{FED}}$. So how does the welfare change when CERT starts paying money to the benign identifier? To answer this question, we characterize the expected user loss and industry loss expressions. Based on those characterization, we state the following

Proposition 6.1 *For $\gamma < \overline{\gamma}_{FED}$, the federally funded social planner outperforms both the CERT-type mechanism and the regulated market-based mechanism along both the metrics – the total user loss and the total industry loss.*

This is a very interesting result. Essentially, we argue that the CERT-type mechanism will be better off if it starts paying out some monetary benefits to the benign identifier, especially if the probability of the vulnerability being reported voluntarily is low. By incentivizing the benign identifier, a federally funded social planner imposes a negative externality on the attacker. Overall, this leads to a better social outcome.

If monetary payment is difficult to implement, one can argue that even non monetary benefits might generate similar results. Therefore, via some non monetary benefit (e.g., due recognition of the identifier), CERT would be able to improve social welfare.

7 Conclusion

In conclusion, we analyze the implication of this market-based mechanism relative to other mechanisms and show the following:

- Contrary to market efficiency arguments, a monopolistic market-maker in an unregulated framework deteriorates the user welfare to the extent that it is almost always worse than even doing nothing. This is because in an unregulated market, the monopolistic market-maker always has an incentive to "leak" any vulnerability it receives from the benign identifier without proper safeguards. This serves to reduce the overall welfare.
- When users voluntarily provide vulnerability information, the market-based mechanism does not perform as well as the CERT-type mechanism even when it is regulated. When voluntary disclosure is low, encouraging a market-based mechanism but with regulation is a good idea.
- Finally, the best mechanism to implement is to let "CERT" fund vulnerability discovery.

While our results have interesting implications, our analysis is not without limitations. For tractability reasons, we use specific functional forms. One future direction would be to use more general functional forms. In addition, we also assume that attacks on software users occur instantaneously which can be generalized as well. We can also consider duopoly market structure and see whether it performs better than CERT. Another interesting extension of this paper would be to empirically validate our model.

References

Arora, A., Caulkins, J.P. & Telang, R. (2003). Provision of Software Quality in the Presence of Patching Technology, Carnegie Mellon University, working paper.

Banker, R., Davis, G. & Slaughter, S. (1998). Software Development Practices, Software Engineering Complexities, and Software Maintenance. *Management Science*, **44**, 433–450.

CERT (2003). CERT/CC Statistics 1988-2003, `http://www.cert.org/stats/`.

Dasgupta, P. & Stiglitz, J. (1980). Uncertainty, Industrial Structure, and the Speed of R&D. *Bell Journal of Economics*, **11**, 1–8.

Du, W. & Mathur, A. (1998). Categorization of Software Errors that led to Security Breaches. In *21st National Information Systems Security Conference, Crystal City, VA*, 392–407.

e Week (2003). CERT, Feds Consider New Reporting Process, `http://www.eweek.com/article2/0,3959,970574,00.asp`.

Gordon, L.A. & Loeb, M.P. (2002). The Economics of Information Security Investment. *ACM Transactions on Information and System Security*, **5**.

Krsul, I., Spafford, E. & Tripunitara, M. (1998). Computer Vulnerability Analysis. Tech. rep., Department of Computer Science, Purdue University, `citeseer.nj.nec.com/krsul98computer.html`.

Varian, H.R. (2000). Managing Online Security Risks. *The New York Times*, `http://www.nytimes.com/library/financial/columns/060100econ-scene.html`.