

Economic Analysis of the Market for Software Vulnerability Disclosure

Karthik Kannan
Purdue University
kkarthik@purdue.edu

Rahul Telang
Carnegie Mellon University
rtelang@andrew.cmu.edu

Hao Xu
Carnegie Mellon University
xhao@andrew.cmu.edu

October 1, 2003

Abstract

Software vulnerability identification and their disclosure has been 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 and users of the software. After verifying a reported vulnerability, and obtaining the remediation in the form of a patch from the software vendor, the infomediary – CERT – sends out a public “advisory” to inform software users about it. In the CERT-type mechanism, reporting vulnerabilities is voluntary with no explicit monetary gains to benign identifiers.

Of late, firms such as iDefense have been proposing a different market-based mechanism. In this market-based mechanism, the infomediary rewards identifiers for each vulnerability disclosed to it. The infomediary then shares this information with its clients who are users of this software. Using this information, clients can protect themselves against attacks that exploit those specific vulnerabilities. The key issue addressed in this paper is whether movement towards such a market-based mechanism for vulnerabilities leads to a better social outcome? We study this problem by characterizing the behavior of software users benign and malign identifiers (or hackers).

1 Introduction

Software vulnerability identification and their disclosure has been a critical area of concern for policy makers. Traditionally, organizations like Computer Emergency Response Team (CERT), have been acting as centralized repositories where identifiers report vulnerabilities. Then, they contact software vendors for the appropriate patches. They also disclose these vulnerabilities publicly after an optimal time. This process allows software users to secure their systems against known and disclosed vulnera-

bilities. In this traditional mechanism, reporting vulnerabilities is voluntary with no explicit monetary gains to identifiers. Of late, commercial firms have been proposing a different market-based mechanism to induce identifiers into providing vulnerability information to them. Firms then share this information with their clients and provide services to secure their clients’ architecture. For example, iDefense (www.iddefense.com) is one such active infomediary in this market.

The key issue is whether movement towards such a market-based mechanism for vulnerabilities leads to a better social outcome? The answer to this question is not obvious. On one hand, incentives to find and disclose vulnerabilities may lead to benign identifiers investing more effort and time in finding vulnerabilities leading to better security. On the other hand, this mechanism may lead to a market with higher levels of vulnerabilities. Although software users which are clients of infomediarities such as iDefense are protected, the rest of the market suffers. The main goal of this paper is to analyze the welfare implications of these existing infomediary mechanism designs. In addition to this, we compare these mechanism designs to two other designs we propose in this paper.

We study this emerging and important research issue by modelling the optimal behavior of all players involved in this market (identifiers, software users and infomediary). Specifically, we model the competition between a benign identifier and malign identifier (or *hackers*). Note that a hacker benefits from vulnerability search and disclosure in a way different from benign identifiers. This, in turn, could lead to a higher or lower hacking incidents affecting social welfare. The analysis provides us with a deeper understanding of the dynamics of this market to apply the right kind of policy interventions.

The paper is organized as follows. In section 2, we review the literature most relevant to this topic. Following that in section 3, we model different mechanism designs for the infomediary. In section 3.6, we compare these mechanism designs.

2 Literature Review

Most prior work in the software vulnerability and information security area has focused on the technical aspects of the problem. But in this section, we restrict our attention to papers that address “non”-technical issues. For example in [8], each vulnerability is categorized based on the policies that get violated when that vulnerability is exploited. Their analysis contributes to the understanding of the steps needed to eradicate these vulnerabilities. [3] take this classification one step further. They categorize and analyze software errors that led to security breaches. Based on their classification schemes, they also develop testing techniques in [4] that can identify security errors.

When such techniques are incorporated in the software development processes, a few papers [7, 2] argue that software quality improves and the software has fewer bugs. Although these methods and processes are useful in improving software quality, 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. Gordon and Loeb [6] develop an economic model for information security investment decisions. They claim that the optimal information security spending does not always increase with the expected loss from attacks and that the optimal security spending has to be far less than the expected loss from attacks. They provide intuitions for this counter-intuitive result and validate their claims using empirical data. Similarly, Arora *et al.* [1] 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. They compare the decision process of a social-welfare maximizing monopolistic vendor to that of a profit-maximizing monopolistic vendor. Interestingly, they observe that the profit-maximizing vendor delivers a product that has fewer bugs than a social-welfare maximizing vendor. However, the profit-maximizing vendor is less willing to patch its software than its social-welfare maximizing counterpart.

To our knowledge, no prior work has addressed issues related to the current imbroglio. Practitioners in different capacities have been proposing different legal/economic frameworks for software vulnerability disclosure ([9, 5, 11]). A few researchers have suggested other mechanisms as well. For example in a New York Time article, Varian [10] suggests that information security can be improved by first assigning legal liability. Along with a legal framework, he argues that an insurance framework can provide the correct market-based incentive structure. In this current scenario, policy-makers are left with little guidance in understanding the implications of these proposed frameworks. Before policy-makers ex-

plore these proposed frameworks, they need a better understanding of the implications of existing frameworks and this is the contribution of our paper. Our paper employs game-theoretic models to provide insights into the welfare-metrics of existing disclosure frameworks.

3 Model

There are four types of participants in this marketplace – the information intermediary, benign identifier, malign identifier and software users. We are interested in comparing the welfare-effects when mechanism adopted by infomediary changes. Specifically, we compare the welfare-effects of the following mechanisms:

- **Market-Based Mechanism:** This is the model implemented by a market-based firm like iDefense which is interested in maximizing profits. It purchases vulnerability information from a benign identifier and secures the architecture for their clients. Its clientele are software users who pay a fee for obtaining this information.
- **CERT-Type Mechanism:** As the name suggests, this models the traditional framework employed for vulnerability disclosure. In this mechanism, no monetary benefits are provided to the identifier. Also, no subscription fee is charged to software users.
- **Consortium Mechanism:** In this case, the infomediary is maximizing the surplus of all its clients. The planner wants to include optimal number of users so that it can cover its costs.
- **Federally Funded Social Planner:** This is reference model against which we compare all other mechanisms. The objective of this mechanism is to maximize social welfare generated. The mechanism is such that a federally funded agency sponsors the vulnerability disclosure. But unlike in the market-based mechanism or the consortium mechanism, no fee is charged to the software users. Vulnerability information is provided to all software users.

We model each of these mechanisms as a two-staged game. In the first stage, the infomediary moves and offers its pricing strategy. In the second stage, all other participants react. The infomediary sets p_b and p_s in the first stage. p_b is the price that the infomediary pays for each vulnerability reported. p_s is the price it charges each of its clients – software users. These prices determine the number of subscriber (and hence the market share) to this service, number of vulnerabilities reported and the probability of attacks. Essentially, they dictate the social welfare generated for each (p_b, p_s) pair. The expression for

social welfare in terms of p_b and p_s is computed in section 3.1. Having done that, we compute the optimal p_b, p_s pair and its corresponding social welfare for each mechanism in each of the following subsections. In the last subsection, we compare the welfare effects of the four mechanisms.

3.1 Modeling Software Users, The Benign Identifier and The Hacker

In this section, we model the reactions of software users, the benign identifier and the hacker to any (p_b, p_s) pair set by the infomediary. We begin by labelling the parameters of interest to us. Let N represent the number of users subscribed to the infomediary's service. This is dependent on the one-time subscription fee p_s charged by the infomediary. Recall that the infomediary acquires the vulnerability information by paying a reward of p_b for each vulnerability reported. This reward incentivizes a benign identifier to exert effort, discover vulnerabilities and report them. Without loss of generality, we assume that there is one vulnerability that can be discovered by either benign identifier or malign identifier with some probability.

Let K_{reported} be the probability that the benign identifier discovers and reports the vulnerability to the infomediary. After obtaining the vulnerability information, the infomediary notifies¹ its clients so that they can protect their systems against potential future attacks. Let the probability that the attack is prevented be $K_{\text{prevented}}$. Thus $K_{\text{prevented}}$ corresponds to the probability that the vulnerability reported by the benign identifier is discovered later by the hacker. In such a case, the hacker can exploit the vulnerability to attack those users that are not subscribed to the infomediary's service. Sometimes the hacker may discover the vulnerability first. Let K_{hacker} be the probability that the vulnerability is first discovered by the hacker. In this case, the hacker exploits the vulnerability to attack all users. All these probabilities – K_{reported} , $K_{\text{prevented}}$, and K_{hacker} – are dependent on the effort level exerted by the benign identifier and the hacker which are in turn, driven by p_s and p_b set by the infomediary. Our objective in this section, is to express these probabilities and N as functions of p_s and p_b . We intend to use these expressions to compute the optimal p_b and p_s under each mechanism.

3.1.1 Software Users

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 uniformly between $[0, \bar{\theta}]$. Any software user i of type θ_i is assumed to

¹The infomediary may also provide value added services such as delivering the patch for the vulnerability, filters to protect against attacks that exploit the vulnerability etc.

incur a loss of θ_i^2 when the vulnerability is exploited. The non-linear choice of the loss function reflects the empirical observations quite well – many users suffer smaller losses while a few users suffer huge losses. The software users have an option of preventing attacks on their systems by subscribing to the infomediary's service. Any user i , whose expected profit from subscribing

$$\Pi_{\text{user}} = \theta_i^2 K_{\text{prevented}} - p_s > 0 \quad (1)$$

subscribes to the service. In this expression, the first term corresponds to the loss prevented by subscribing to the service. The second term corresponds to the payment made to the infomediary. Clearly only those software users whose θ_i satisfies the following condition subscribe to the service:

$$\theta_i > \sqrt{\frac{p_s}{K_{\text{prevented}}}} \quad (2)$$

Since θ is assumed to be uniformly distributed between $[0, \bar{\theta}]$, the number of clients subscribed to the infomediary's service is:

$$N = 1 - \sqrt{\frac{p_s}{K_p}} \frac{1}{\bar{\theta}} \quad (3)$$

Consider the CERT mechanism where software users are not charged any price at all i.e., $p_s = 0$. In such a case, $N = 1$ which implies all users are provided with vulnerability information.

3.1.2 Competition Between the Benign Identifier and the Hacker

Given p_b and p_s , how do benign users and hackers change their effort level? In short, we are interested in obtaining the probabilities – K_{reported} , $K_{\text{prevented}}$ and K_{hacker} – as functions of p_b and p_s . As a first step, we express these probability values in terms of effort levels of the benign identifier and the hacker. Then in the second step, we compute the optimal effort level exerted by them based on their respective profits.

The competition between the benign identifier and the hacker is modelled by considering the time-frame for the software's life cycle as T . Within this period, we assume that the probability that a player – the benign identifier or the hacker – discovers a vulnerability is distributed uniformly. Note that we are simply characterizing the probability of discovering the vulnerability and not the probability conditioned on a player discovering it first. Let γ , an exogenous parameter, correspond to a baseline probability such that each player discovers the vulnerability before the end of the time period without exerting any effort. Therefore, given our distributional assumption, $\frac{\gamma}{T}$ is the probability that the vulnerability is discovered by a player at each instant without exerting any

effort. In short this is the *pdf* for vulnerability discovery at some time t . Players can alter γ and hence this *pdf* by exerting effort. For simplicity we assume that when a user exerts an effort α , the probability of discovering the vulnerability $\alpha + \gamma^2$. Similarly, for an effort level β , the probability increases to $\beta + \gamma$.

Obviously, both benign user and hacker would invest in efforts α and β in an optimal manner. In short, they are determined by the non cooperative Nash equilibrium that emerges due to competition between them. These parameters α and β – are assumed to be set for the entire duration, T , and cannot be modified during the game. Because of their efforts, the different probability values are computed as follows:

- The probability that the vulnerability is reported – K_{reported} – corresponds to the probability that the vulnerability is first discovered by the benign identifier and reported to the infomediary.

$$K_{\text{reported}} = \int_0^T \frac{\text{Probability}(\text{benign} = t)}{\text{Probability}(\overline{\text{hacker}} < t)} dt \quad (4)$$

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

$$\begin{aligned} K_{\text{reported}} &= \int_0^T \frac{\alpha + \gamma}{T} \left(1 - \frac{(\beta + \gamma)t}{T}\right) dt \\ &= (\alpha + \gamma) \left(1 - \frac{(\beta + \gamma)}{2}\right) \end{aligned} \quad (5)$$

- $K_{\text{prevented}}$ is the probability with which an attack is prevented because a benign user identified it before the hacker. Therefore, this is also the value of subscribing to the infomediary's service. Therefore,

$$\begin{aligned} K_{\text{prevented}} &= \int_0^T \frac{\text{Probability}(\text{hacker} = t)}{\text{Probability}(\text{benign} < t)} dt \\ &= (\alpha + \gamma) \frac{(\beta + \gamma)}{2} \end{aligned} \quad (6)$$

- Finally, the probability that the vulnerability is first

discovered by the hacker – K_{hacker} – is

$$\begin{aligned} K_{\text{hacker}} &= \int_0^T \frac{\text{Probability}(\text{hacker} = t)}{\text{Probability}(\overline{\text{benign}} < t)} dt \\ &= (\beta + \gamma) \left(1 - \frac{(\alpha + \gamma)}{2}\right) \end{aligned} \quad (7)$$

Having defined these probabilities and their properties, we can now characterize the optimal efforts exerted by the benign identifier and the hacker. For this, we consider their respective expected profit functions. Recall that the effort exerted by the benign identifier increases its probability of finding the vulnerability to $\alpha + \gamma$. This effort pays p_b , if the benign identifier discovers the vulnerability before the hacker. Since K_{reported} is the probability that the benign identifier discovers the vulnerability first, the expected revenue for the benign identifier is given by $p_b K_{\text{reported}}$. Corresponding to its effort, the benign identifier's cost is $C(\alpha)$, a function of α . Mathematically, the expected profit for the benign identifier is:

$$\Pi_b = K_{\text{reported}} p_b - C(\alpha)$$

For obtaining an interior optimal solution, we require that Π_b be concave with respect to α . Since the revenue increases linearly with α , any convex cost function will suffice. Since it is commonly used, we use the quadratic function, $C(\beta) = M\alpha^2$ where M is an exogenous constant parameter which we use for scaling purpose such that K_{reported} , $K_{\text{prevented}}$, and K_{hacker} are bounded above by one. Substituting for $C(\alpha)$ and K_{reported} ,

$$\Pi_b = (\alpha + \gamma) \left(1 - \frac{(\beta + \gamma)}{2}\right) p_b - M \alpha^2 \quad (8)$$

Next, let us consider the hacker's expected profit. The hacker 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 clientele³. We assume that if the hacker is successful in attacking a user of type θ_i , he gains a profit of θ_i . Note that the functional form of the hacker's profit function is intentionally made to be different from the loss for the user – θ_i^2 . The hacker's cost is $C(\beta)$. Therefore,

$$\begin{aligned} \Pi_h &= K_{\text{hacker}} \left(\int_0^{\bar{\theta}} \theta \frac{1}{\theta} d\theta \right) \\ &+ K_{\text{prevented}} \left(\int_0^{(\bar{\theta} - N\bar{\theta})} \theta \frac{1}{\theta} d\theta \right) - C(\beta) \end{aligned}$$

²We found that the results are robust to the functional form assumptions made. We obtained a similar set of results when using a multiplicative functional form.

³It is trivial to show that hacker never finds it optimal to sell the vulnerability.

In the first term, K_{hacker} corresponds to the probability that the hacker discovers the vulnerability before the benign identifier and attacks all the users. The term inside the integral is the expected profit that the hacker obtains from attacking all the users. Similarly in the second term, $K_{\text{prevented}}$ corresponds to the probability that the hacker discovers the vulnerability after the benign identifier. The integral in the second term corresponds to the expected profit that the hacker can gain by attacking users that are not part of the infomediary's clientele. The last term corresponds to the cost of exerting effort. Substituting for K_{hacker} and $K_{\text{prevented}}$, and simplifying the expression, we have

$$\begin{aligned} \Pi_h = & (\beta + \gamma) \left(1 - \frac{(\alpha + \gamma)}{2}\right) \frac{\bar{\theta}}{2} \\ & + \frac{(\beta + \gamma)(\alpha + \gamma) \bar{\theta} (1 - N)^2}{2} - M \beta^2 \end{aligned}$$

Substituting for N from equation 3, we have the expected profit for the hacker as

$$\begin{aligned} \Pi_h = & (\beta + \gamma) \left(1 - \frac{(\alpha + \gamma)}{2}\right) \frac{\bar{\theta}}{2} \\ & + \frac{(\beta + \gamma)(\alpha + \gamma)}{2} \frac{p_s}{K_{\text{prevented}} 2\bar{\theta}} - M \beta^2 \end{aligned}$$

We can simplify this expected profit function further by substituting for $K_{\text{prevented}}$. Therefore,

$$\Pi_h = (\beta + \gamma) \left(1 - \frac{(\alpha + \gamma)}{2}\right) \frac{\bar{\theta}^2}{2} + \frac{p_s}{2\bar{\theta}} - M \beta^2 \quad (9)$$

These expected profit expressions are used to determine the optimal values for α and β . To obtain the optimal level of effort for the benign identifier – α – we differentiate the benign identifier's expected profit expression i.e., equation 8, with respect to α and equate it to zero. Thus,

$$\alpha^* = \left(1 - \frac{\beta + \gamma}{2}\right) \frac{p_b}{2M} \quad (10)$$

Similarly to obtain the optimal effort level for the hacker – β – we differentiate the hacker's expected profit expression i.e., equation 9, with respect to β and set it to zero.

$$\beta^* = \left(1 - \frac{\alpha + \gamma}{2}\right) \frac{\bar{\theta}}{4M} \quad (11)$$

Solving the simultaneous equations – equation 10 and equation 11, we get:

$$\alpha^* = \frac{(8M - \bar{\theta}) p_b (2 - \gamma)}{32M^2 - p_b \bar{\theta}} \quad (12)$$

$$\beta^* = \frac{(2 - \gamma)(4M - p_b) \bar{\theta}}{32M^2 - p_b \bar{\theta}} \quad (13)$$

As we mentioned before, since $\alpha + \gamma$ and $\beta + \gamma$ are also probabilities and should be bound from above, these expressions are valid only for $M > M_{th} = \frac{(2 - \gamma)\bar{\theta}^2 + \sqrt{\bar{\theta}^3(2 - \gamma)^2 - 2 - 2\gamma}}{8(1 - \gamma)}$. Therefore, we can scale M such that this inequality is always satisfied.

For $M > M_{th}$, we observe the following properties in these equations: a) Both these expressions are independent of p_s , the one-time fee charged by the infomediary. b) As p_b increases, α increases but β decreases. This suggests that the effort exerted by the benign identifier increases with p_b . But this, in turn, imposes a negative externality on the hacker's incentives and reduce its efforts. c) For a given p_b , both the benign identifier and the hacker have an incentive to increase their efforts as γ decreases. d) Finally, as M increases, i.e., the cost of exerting effort increases, the optimal effort levels – α^* and β^* – decrease as expected.

Using α^* and β^* , we can also compute the probabilities $K_{\text{prevented}}$, K_{reported} , K_{hacker} ⁴. From these equations, we observe the following properties to be valid:

- As p_h , the incentive to disclose vulnerability increases, K_{reported} , the number of vulnerabilities reported increases.
- Similarly, higher p_h increases $K_{\text{prevented}}$, the number of vulnerabilities prevented also increases.
- Finally, K_{hacker} decreases with p_h .

3.1.3 Welfare Metrics and Comparative Statics

Given these behavior, we are ready to compute the welfare metrics – Total User Loss and Total Industry Loss – for any given p_b and p_s . The total user loss is specified as:

$$\begin{aligned} UL = & K_{\text{hacker}} \left(\int_0^{\bar{\theta}} \frac{\theta^2}{\bar{\theta}} d\theta \right) \\ & + K_{\text{prevented}} \left(\int_0^{\bar{\theta}(1-N)} \frac{\theta^2}{\bar{\theta}} d\theta \right) + N p_s \quad (14) \end{aligned}$$

The first expression corresponds to the loss incurred when the hacker discovers the vulnerability first. Note that in this case hacker can attack all users. The second expression corresponds to the loss incurred when the hacker discovers the vulnerability after the benign identifier. In this case, the hacker is left to attack only those users who are not part of the infomediary's clientele. The last term corresponds to the total payment made by all subscribing

⁴Since these probabilities involve complex expressions, we do not provide them. However, they are available upon request.

users to the infomediary. Substituting for different parameters, one can compute the total user loss as a function of p_b and p_s . Since, p_s and p_b are somewhat tedious, this loss function is tedious as well and we do not report it here. We perform numerical analysis.

Similarly, we compute the total industry loss as the loss incurred by the users in addition to the profit/loss incurred by the infomediary. Equation 14 corresponds to the loss incurred by the users. This combined with the infomediary's profit equates to

$$IL = K_{\text{hacker}} \left(\int_0^{\bar{\theta}} \frac{\theta^2}{\bar{\theta}} d\theta \right) + K_{\text{prevented}} \left(\int_0^{\bar{\theta}(1-N)} \frac{\theta^2}{\bar{\theta}} d\theta \right) + K_{\text{reported}} p_b \quad (15)$$

When we compute the industry profits, the term $N p_s$ which appears in equation 14 does not appear in equation 15. This is because $N p_s$ is the transfer of rent between the users and the infomediary. Thus, the only remaining term is the expected payment made by the infomediary for vulnerability disclosure and it appears in equation 15. Substituting for different parameters, one can compute the total industry loss as a function of p_b and p_s .

3.2 Market-based Mechanism

A classic example of this type of infomediary is iDefense (www.iDefense.com). It purchases vulnerability information and notifies its clients about the vulnerability. Recall that the infomediary charges each of its clients a price of p_s and pays p_b for every vulnerability reported. These variables are outcome of the expected profit maximization function given by

$$\max_{p_s, p_b} N p_s - K_{\text{reported}} p_b \quad (16)$$

The first term in the expression corresponds to the revenue that the infomediary generates from charging its clients p_s . The second term is the cost it incurs to pay for each vulnerability reported.

We can substitute N from equation 3. Therefore, the objective function is:

$$\max_{p_s, p_b} \left(1 - \sqrt{\frac{p_s}{K_p} \frac{1}{\bar{\theta}}} \right) p_s - K_{\text{reported}} p_b \quad (17)$$

Taking the first order derivative with respect to p_s and setting it to zero give us optimal p_s . Therefore,

$$p_s^* = \frac{4 K_{\text{prevented}} \bar{\theta}^2}{9}$$

We set $p_s = p_s^*$ in equation 17, differentiate the expres-

sion with respect to p_b and equate it to zero to obtain the optimal p_b^5 .

Given p_s and p_b , we can calculate the welfare metrics for the market mechanism – the total user loss, UL_{MKT} , and the total industry loss, IL_{MKT} . In the subsequent section we show the expression for both these measures.

3.3 CERT-type Mechanism

In this mechanism, $p_b = 0$ i.e., no monetary incentives are provided to the identifier, and $p_s = 0$ i.e., users are not charged any subscription fees for vulnerability notification. All users are notified about the vulnerability. By virtue of our derivation in section 3.1, when $p_b = 0$, $\alpha = 0$ i.e., the benign identifier does not exert any effort at all. But, vulnerabilities are still discovered by a benign identifier with a probability of γ . By assumption, the benign identifier always reports the vulnerability to the infomediary. On the other hand, hacker invests an optimal β to exploit the vulnerabilities. Optimal β is found by substituting $p_b = 0$ in equation 13. Therefore,

$$\beta^* = \frac{(2 - \gamma)\bar{\theta}}{8M} \quad (18)$$

Again we can calculate the welfare metrics for CERT mechanism. But in this case, both losses are equal to one another since there is no transfer of payment. For $p_b = 0$ and $p_s = 0$, the losses are given by

$$UL_{\text{CERT}} = IL_{\text{CERT}} = \frac{1}{48M} (2 - \gamma)((2 - \gamma)\bar{\theta} + 8M\gamma)\bar{\theta}^2 \quad (19)$$

3.4 Consortium Mechanism

Although currently, no equivalent framework exists, one can imagine Information Sharing and Analysis Centers (ISAC)⁶ to execute such a mechanism. In this mechanism, the infomediary is assumed to maximize the welfare generated for all its clients. It does so by charging each client a one-time fee of p_s that is sufficient enough to pay for the vulnerabilities reported:

$$N p_s = K_{\text{reported}} p_b \quad (20)$$

The left hand side of this expression is the income to the infomediary and the right hand side is the expected pay-

⁵Here again, the expression is tedious and it is available upon request

⁶The US federal government, under Presidential Decision Directive NSC-63, has encouraged the establishment of industry based Information Sharing and Analysis Centers (ISACs) to promote the disclosure and sharing of security information among firms. Currently, these ISACs are focused on gathering, analyzing and sharing information related to actual, as well as unsuccessful attempts at, security breaches. We envision their role to be broader for this model.

ment made by the infomediary for the vulnerability discovery. Substituting for N , we have

$$\left(1 - \sqrt{\frac{p_s}{K_{\text{prevented}} \frac{1}{\bar{\theta}}}}\right) p_s = K_{\text{reported}} p_b$$

This is a constraint to the following objective function optimized by the infomediary:

$$\max_{p_b} \left(\int_{\bar{\theta}(1-N)}^{\bar{\theta}} \theta^2 \frac{1}{\bar{\theta}} d\theta \right) K_{\text{prevented}} - N p_s$$

The first term corresponds to the loss prevented by this framework whereas the second term corresponds to the cost incurred by all clients.

We express this constrained optimization function used by the infomediary in the following manner:

$$\max_{p_b, p_s, L} \left(\int_{\bar{\theta}(1-N)}^{\bar{\theta}} \theta^2 \frac{1}{\bar{\theta}} d\theta \right) K_{\text{prevented}} - N p_s + L(K_{\text{reported}} p_b - N p_s)$$

where L is the Lagrange variable. We solve this function using Kuhn-Tucker method to obtain p_s and p_b , and then the corresponding welfare metrics. For simplicity reason, we do not list the results here. We will compare it to other mechanisms later.

Note that similar to the CERT-framework, the industry loss IL_{CON} and the loss incurred by the users UL_{CON} are identical i.e., $UL_{\text{CON}} = IL_{\text{CON}}$. This is so because the transfer of payment is equal to the loss incurred by the infomediary. Since p_s and p_b are analytically intractable, one can compute these losses numerically.

3.5 Federally-Funded Social Planner

This subsection deals with the case when the infomediary pays for vulnerability disclosure but charges nothing to the users i.e., $p_s = 0$. The price that it is willing to pay for the vulnerabilities is a solution to the following optimization function

$$\max_{p_b} \left(\int_0^{\bar{\theta}} \theta^2 \frac{1}{\bar{\theta}} d\theta \right) K_{\text{prevented}} - K_{\text{reported}} p_b \quad (21)$$

Solving this optimization function, we obtain the optimal p_b^* . In this expression, p_b decreases with an increase in γ . $p_b > 0$ only for $\gamma < \frac{\bar{\theta}^5}{48 M^2 - 4M \bar{\theta} + \bar{\theta}^5}$. Beyond this, the performance will be similar to that of CERT. Corresponding to these prices, we evaluate the total user loss, UL_{FED} , and total industry loss, IL_{FED} .

3.6 Comparison

We use results from the previous subsections to compare the four mechanisms. We study the sensitivity of the total user loss to variations in each of γ , M and $\bar{\theta}$. For simplicity of exposition, we will refer to the total user loss as TUL for the rest of this section. Our results are:

- Keeping M and γ constant, as $\bar{\theta}$ increases, i.e., as the maximum loss suffered by the user increases, TUL , increases under all four cases. This is observed in Figure 1.
- Figure 2 shows the sensitivity of TUL to M for a set value of $\gamma = 0$ and $\bar{\theta} = 6$. We observe that as M increases, both TUL decrease. Intuitively as the cost of discovering vulnerabilities increases, the optimal effort level exerted by the hacker decreases which in turn, decreases the loss suffered by users.
- Figure 3 shows the variations of TUL due to γ . Recall that γ is the random independent probability that the vulnerability is discovered without any effort. As γ increases, we observe the user loss to increase.
- Also note that $\gamma = 0$, CERT-type mechanism performs worse than a Market-based Mechanism.
- In contrast at higher values of γ , we observe the reverse i.e., CERT-type mechanism performs better than a Market-based Mechanism.
- Federally Funded Social welfare maximizing mechanism always performs better than all other mechanisms.

The results were found to be exactly similar when studying the total industry loss also.

4 Conclusion

In conclusion, we have studied an important real-world problem related to software vulnerability disclosure. Specifically, we use game-theory to compare the welfare-effects of different mechanisms. Based on our analysis, we observe that a Federally-Funded Social Planner performs better than all other mechanisms. We also find that under certain conditions, the performance of market-based mechanism is better than the CERT-type one and vice-versa. This understanding is critical especially given the current scenario where practitioners in different capacities are proposing new frameworks. We intend to extend this work by comparing other disclosure frameworks proposed by different practitioners.




Figure 1: Total User Loss for different Mechanism Framework as $\bar{\theta}$ changes. $M = 15$ and $\gamma = 0$ for this analysis.




Figure 2: Total User Loss for different Mechanism Framework as M changes. $\bar{\theta} = 6$ and $\gamma = 0$ for this analysis.

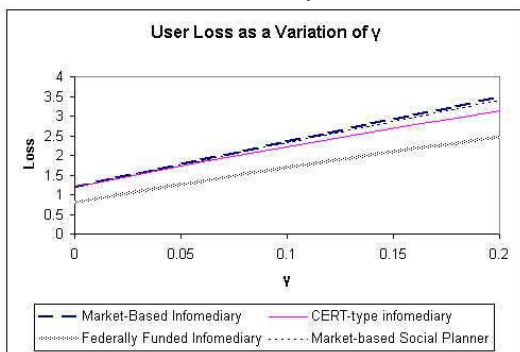


Figure 3: Total User Loss for different Mechanism Framework as γ changes. $\bar{\theta} = 6$ and $M = 15$ for this analysis.

References

- [1] Ashish Arora, Jonathan P. Caulkins, and Rahul Telang. Provision of software quality in the presence of patching technology. Carnegie Mellon University, working paper, 2003.
- [2] Rajiv Banker, G. Davis, and Sandra Slaughter. Software development practices, software engineering complexities, and software maintenance. *Management Science*, 44(4):433–450, 1998.
- [3] Wenliang Du and A.P. Mathur. Categorization of software errors that led to security breaches. In *21st National Information Systems Security Conference, Crystal City, VA*, pages 392–407, 1998.
- [4] Wenliang Du and A.P. Mathur. Vulnerability testing of software system using fault injection. Technical report, Department of Computer Science, Purdue University, 1998. Reference: Coast TR 98-02.
- [5] e Week. Cert, feds consider new reporting process. <http://www.eweek.com/article2/0,3959,970574,00.asp>, 2003.
- [6] Lawrence A. Gordon and Martin P. Loeb. The economics of information security investment. *ACM Transactions on Information and System Security*, 5(4), 2002.
- [7] Mayuram S. Krishnan, Charlie H. Kriebel, Sundar Kekre, and Tridas Mukhopadhyay. An empirical analysis of productivity and quality in software products. *Management Science*, 46(6):745–59, 2000.
- [8] Ivan Krsul, Eugene Spafford, and Mahesh Tripunitara. Computer vulnerability analysis. Technical report, Department of Computer Science, Purdue University, May 1998. citeseer.nj.nec.com/krsul98computer.html.
- [9] Security-Focus. Security research exemption to dmca considered. <http://www.securityfocus.com/news/4729>, 2003.
- [10] Hal R. Varian. Managing online security risks. *The New York Times*, 2000. <http://www.nytimes.com/library/financial/columns/060100econ-scene.html>.
- [11] ZD-Net. Trusted computing comes with a warning. <http://www.zdnet.com.au/newstech/security/story/0,2000048600,20273805,00.htm>, 2003.