Bug Bounty Programs, Security Investment and Law Enforcement: A Security Game Perspective

Jiali Zhou[*] Kai-Lung Hui[†]

May 15, 2019

Abstract

Bug bounty programs are gaining popularity, but practitioners have not agreed on their effectiveness. We use a stylized model to analyze the economic trade-offs in bug bounty programs. Our analysis provides six main insights: 1) Unless a firm needs to pay an excessive bounty reward to entice a strategic hacker to participate, it is always beneficial for the firm to launch a bug bounty program. (2) The firm enjoys two benefits from a bug bounty program: attack diversion and protection delegation. (3) The firm optimally retains in-house protection; the level of in-house protection and the bounty prize depend on the balance of two incentives: bounty-payment squeezing and protection free-riding. (4) Law enforcement affects private protection only when the level of enforcement is relatively small. In this case, strengthening law enforcement increases the company's payoff but could variously increase or decrease the strategic hacker's payoff. Excessive enforcement may make a system less secure. (5) When the exogenous threat is significant, the company prefers a more capable participant. But when the exogenous threat is small, the company prefers a complementary participant. (6) Bug bounty programs need not provide better security protection.

1 Introduction

Every organization has valuable assets to secure. This need becomes critical and challenging in the cyberspace. In general, the security of organizational assets depends on two factors. One is an organization's in-house protection measures such as security soft(hard)ware that detects and prevents intrusions, security audits that reduce system bugs, and cybersecurity awareness and training. The second is public deterrence measures such as cyber-laws which impose punishment or polices which increase conviction rates. Existing studies on how to optimize each of these measures

[*] School of Business and Management, Hong Kong University of Science and Technology. jzhoubf@connect.ust.hk

[†] School of Business and Management, Hong Kong University of Science and Technology. klhui@ust.hk

are extensive, and both sets of measures improve security essentially by reducing hacker's gain from the attack.

As new protection measures emerge, organizations have to decide whether and how to add new measures into their protection. The bug bounty program is an emerging measure that taps into the "wisdom of crowd" to improve system security (Maillart et al. 2017). It refers to a setting where a software vendor pays an external party a bug bounty reward for disclosing the details of a security-related vulnerability (Egelman et al. 2013). Many large organizations, including Google, Tencent, and Facebook have been offering bug bounty programs. In 2017 alone, Google paid $2.9 million in bug bounty programs, with $112,500 being the biggest single reward (Keller 2018). In 2016, United States Department of Defense launched a bug bounty program called "Hack the Pentagon". It recently expanded this program with a ceiling reward of $34 million (Fbo.gov 2018).

Notwithstanding its popularity, many organizations are reluctant to use bug bounty programs. For example, Oracle's chief security officer, Mary Ann Davidson, claimed that bug bounties are uneconomic and "the new boy brand", and that "we find 87% of security vulnerabilities ourselves, security researchers find about 3% and the rest are found by customers…on a strictly economic basis, why would I throw a lot of money at 3% of the problem when I could spend that money on better prevention like, oh, hiring another employee to do ethical hacking." (Oracle.com 2015). Similarly, Microsoft once claimed that "we dont think paying a per-vuln bounty is the best way" (Fisher 2010) although it later offered a bug bounty program in 2013.

In this paper, we formally examine the economic trade-offs in using a bug bounty program. Our analysis encompasses a firm owning a valuable and vulnerable system. To protect the system from cyber-attacks, the firm is considering whether to invest in a bug bounty program. With this setting, we study the nature of a bug bounty program by answering the following questions: (1) Should the firm offer a bug bounty program? If so, how should the reward be set and what expected benefits does such a program bring? (2) How does a bug bounty program affect the incentives of the firm to invest in in-house protection? (3) How does a bug bounty program interact with law enforcement? Can it help the firm arrive at better overall protection?

We develop a stylized model to address these questions. In our model, a firms wants to protect a vulnerable system and a strategic hacker wants to identify the vulnerability for economic benefits. There are some naive hackers who always want to attack the system. The firm can focus on in-house protection or offer a bug bounty program. Without a bug bounty program, the strategic hacker can benefit from directly exploiting the vulnerability, e.g., by stealing the firm's customer data and selling them through the black market or mis-using the customers' bank accounts. With a bug bounty program, the strategic hacker can exchange the vulnerability for a bounty reward.

Our analysis shows that the firm should offer a bug bounty program as long as the value that a strategic hacker can obtain from attacking the system is not too high relative to the firm's potential loss. When the bounty program is offered, the firm enjoys two benefits. The first benefit is attack diversion, where the firm can use a cheap price to divert the strategic hacker away from attacking the system. The second benefit is protection delegation, where the firm can engage the strategic hacker in protecting the system against exogenous threats.

When using the bug bounty program, the firm's best strategy involves retaining in-house protection capabilities. The incentive is to squeeze the bug bounty payment. Specifically, how the firm designs the bug bounty program and selects its in-house protection strategy depend on the balance of two incentives, viz. bounty-payment squeezing and protection free-riding. When the exogenous threat is small, the firm tends to set a smaller bounty reward, which is just enough to divert the strategic hacker away from attacking. It will respond to an increase in the strategic hacker's effort by strengthening its in-house protection, so as to reduce the probability of paying the bounty reward, an incentive that we call bounty payment squeezing. When the exogenous threat is significant, the firm tends to offer a bigger bounty reward, the size of which increases with the threat. It will respond to an increase in the strategic hacker's effort by cutting back in-house protection, so as to incentivize more effort from the strategic hacker, an incentive that we call protection free-riding.

We also analyze how law enforcement affects the bug bounty program. We find that law enforcement affects the firm, the strategic hacker, and the bug bounty program only when the level of enforcement is relatively small. When this happens, strengthening law enforcement increases the firm's payoff but could variously increase or decrease the strategic hacker's payoff. Notably, strengthening law enforcement could even generate negative results: The firm may cut back its in-house protection, and the effectiveness of the bug bounty program could reduce, leading to a less secure system.

Finally, we study whether the bug bounty program improves the overall security. Surprisingly, a bug bounty program does not necessarily lead to better security. This unexpected consequence happens when the firm cuts back too much in-house protection, which offsets the security-enhancing effects the bug bounty program brings.

The rest of this paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the model. Section 4 analyzes the equilibrium outcomes and examines our research questions. Section 5 concludes the paper.

2 Related Literature

This study is related to three streams of research. The first stream is the economics of crowdsourcing. As reasoned by Fryer and Simperl (2017), bug bounty program can be conceptually viewed as a crowdsourcing event because it appeals to a defined crowd (security researchers) with a clear goal (identify vulnerabilities) and a defined benefit for both the workers (bounty reward) and the requester (higher security). This literature lies at the intersection of the economics of outsourcing (e.g. Lacity et al. 2009) and the theory of contests and tournaments (e.g. Szymanski 2003, Konrad 2007), and is primarily interested in the optimal design of crowdsourcing contests, such as award structure (e.g. Terwiesch and Xu 2008) , incentive scheme (e.g. Horton and Chilton 2010), participation policies (e.g. Jeppesen and Lakhani 2010, Boudreau et al. 2011), etc.

Despite the similarity, the bug bounty program differs from traditional crowdsourcing settings in three notable ways, making previous analysis less applicable. First, in traditional settings, participants mostly have no bargaining power so they are assumed to be "contract takers", whereas participants in bug bounty programs have substantial bargain power. For example, if an organizer refuses to make a deal or offers too little rewards, it is possible for the bug finders to monetize the vulnerability elsewhere, e.g., by directly attacking the organizer's systems, which in turn poses a threat to the organizer. Second, in previous settings, firms' internal activities are mostly independent of the crowdsourcing program outcomes. By contrast, in a bug bounty program, the firm's in-house protection will interact with the participants' strategies. For instance, higher in-house protection by the firm is likely to decrease the chance for participants to identify the vulnerability and claim the bounty rewards. Third, in traditional contest-type of crowdsourcing programs, quality is the most important measure of participants' submissions. While in a bug bounty program, the uniqueness of the submitted vulnerability is more important.

The second relevant stream is the literature on information security investment. This literature essentially focuses on how to allocate resources to minimize security costs. Following the seminal work of Gordon and Loeb (2002), many studies have discussed how companies optimize security investments under different settings and contexts, such as when there exists system interdependency (e.g. Hausken 2006), when hackers strategically identify poorly protected system (e.g. Cremonini and Nizovtsev 2009, Bandyopadhyay et al. 2014), and when companies share information (Gal-Or and Ghose 2005, Gao et al. 2015). Our main contribution to this literature is analyzing the bug bounty program as an alternative security investment. This is interesting for two reasons: (1) an in-depth understanding of bug bounty program is necessary to guide companies in setting up such a program; (2) compared with previous security investment approaches, bug bounty programs create fundamentally different incentives for the hackers. Previous security investment approaches mostly aim to discourage hackers from identifying system vulnerabilities. Bug bounty programs, on the other hand, encourage hackers to identify vulnerabilities as long as they have incentives to submit the vulnerabilities for bounty rewards.

The last stream of research comprises studies specifically focusing on bug bounty programs. Finifter et al. (2013), Zhao et al. (2015) and Ruohonen and Allodi (2018) provide descriptive analysis of the existing bug bounty programs. They characterize many interesting aspects about these programs, such as their economic efficiency, evolution of the discovered vulnerabilities, and the participants behavioral patterns and trends, etc. Other studies examine the economic trade-offs related to bug bounty programs, such as the relationship between the bounty reward and vulnerability severity (Munaiah and Meneely 2016), the competition between different bug bounty programs (Maillart et al. 2017), and the optimal mechanism to minimize invalid reports and allocate the participants' efforts (Zhao et al. 2017). To our knowledge, systematic investigation of a firm's incentive to launch a bug bounty program does not exist. This is what we contribute in this study.

3 Model

3.1 Setting

Consider a market with one firm who is planning security investment to protect a vulnerable system and one strategic hacker who can benefit from exploiting the vulnerability in this system.[1] For simplicity, assume both the firm and the strategic hacker estimate the system has one vulnerability through some preliminary costless security assessments. The firm can invest effort $q \in [0, 1)$ to identify the vulnerability. Equivalently, q can be interpreted as the probability of finding the vulnerability, so $q \in [0, 1)$. The total effort cost is mq^2, where m is sufficiently large that prevents $q = 1$, i.e., perfect security. If the vulnerability is left unfixed and eventually exploited by others, then the firm will incur a monetary loss of λ.

Similarly, the strategic hacker can invest effort $p \in [0, 1)$ to identify the vulnerability, with cost kp^2. Without a bug bounty program, the strategic hacker can monetize the vulnerability only by attacking the system. We assume the gain in vulnerability exploitation, $\alpha\lambda$, is proportional to the value of the system, where $0 < \alpha < +\infty$. However, vulnerability exploitation is risky. By launching an attack against the firm's system, the strategic hacker faces a probability, d, of being caught and punished with a penalty, f. Both d and f are affected by government enforcement and exogenous to the firm's and strategic hacker's decisions.[2]

[1] Following Kannan and Telang (2005) and Png and Wang (2009), we assume there is only one strategic hacker. We will extend the analysis to N strategic hackers in future work.

[2] A remark: in reality the strategic hacker may also monetize the vulnerability in other ways, e.g. sell the

Table 1: Model Notation

Variable	Description
q	the firm's vulnerability identification effort
p	the strategic hacker's vulnerability identification effort
t	the naive hackers' vulnerability identification effort
m	cost coefficient for the firm's effort
k	cost coefficient for the strategic hacker's effort
λ	the firm's monetary loss from attack
α	coefficient for the strategic hacker's gain from attack
d	probability of being caught if the strategic hacker decides to attack the system
f	expected penalty if the strategic hacker were caught
R	the size of bug bounty reward

With a bug bounty program, the strategic hacker can submit his discovery of the vulnerability to the firm and earn a reward, R, which is set by the firm. The firm can fix the vulnerability as soon as it is discovered. For simplicity, we assume the strategic hacker can exploit the vulnerability and the firm can fix the vulnerability without incurring any cost.

The market also contains some naive hackers who always invest effort $t \in [0, 1)$ in identifying and exploiting the vulnerability in the firm's system. We can interpret t as the probability of naive hackers finding and exploiting the vulnerability. These naive hackers may attack the firm for fun, ideology, or other malicious motives, and they are die-hard attackers who will not be attracted by the bug bounty program.[3] If the firm, the strategic hacker, and naive hackers simultaneously find the vulnerability, we assume the firm will take action first, followed by the strategic hacker and then the naive hackers.[4] We assume the size of naive hackers is large to the extent that their decision is not affected by the firm's and strategic hacker's decisions, i.e., t is exogenous.

Without a bug bounty program, the firm and the strategic hacker simultaneously decide how much efforts to invest in protecting or attacking the system. With a bug bounty program, the firm will announce and commit to a bug bounty reward first before all parties make their effort choices. We focus on subgame-perfect Nash equilibria. Table 1 summarizes all the notations used in this paper.

vulnerability through the black market. But as long as the expected benefit he can get increases with the value of the system, and decreases with the government enforcement, (both of which we believe are reasonable assumptions) the analysis still applies.

[3] We can also put some benign hackers, i.e. who will submit the vulnerability anyway, in the market. It is a straightforward extension but there is no reason to believe it would bring additional insights.

[4] We will change this sequence of action in future work

3.2 Payoffs

3.2.1 No Bug Bounty Program

Without a bug bounty program, the strategic hacker's utility is

$$U_{attack} = p(1-q)(\alpha\lambda - df) - kp^2, \tag{1}$$

where $p(1-q)$ is the probability that the strategic hacker has identified the vulnerability first,[5] $\alpha\lambda$ is the expected gain and df is the expected penalty from exploiting the vulnerability[6].

Following the security investment literature(e.g. Gordon and Loeb 2002), the firm's goal is to maximize the expected net benefit from the security investment. Therefore, given the external threats represented by p and t, the firm's payoff is

$$\Pi_{attack} = -[1-(1-t)(1-p)](1-q)\lambda - mq^2. \tag{2}$$

Here, $[1-(1-t)(1-p)](1-q)$ is the probability that the strategic hacker or naive hackers have identified and exploited the vulnerability.

3.2.2 With a Bug Bounty Program

By submitting the vulnerability to the firm, the strategic hacker's payoff is

$$U_{bounty} = Rp(1-q) - kp^2. \tag{3}$$

He will submit the vulnerability to the firm if and only if the reward exceeds the benefit he would obtain from exploiting it himself, i.e., $U_{bounty} \geq U_{attack}$. If the strategic hacker participates in the bug bounty program, the firm's expected payoff is

$$\Pi_{bounty} = -t(1-p)(1-q)\lambda - p(1-q)R - mq^2. \tag{4}$$

The first term in (4) is the firm's expected loss when naive hackers find the vulnerability earlier than the firm and the strategic hacker. The second term is the expected (bug bounty) payment to the strategic hacker when he finds the vulnerability earlier than others.

[5] Recall if the strategic hacker and naive hackers (but not the firm) find the vulnerability simultaneously, then the strategic hacker will take action before naive hackers.

[6] This is a standard setup in criminology literature (e.g. Becker 1968, Chalfin and McCrary 2017)

4 Analysis

We start with a simple setting where the firm and the strategic hacker have similar capability in identifying the vulnerability, i.e., $m = k$. We relax this assumption in a later section.

4.1 Equilibrium With No Bug Bounty Program

Differentiating the strategic hacker's payoff function in (1), his optimal investment is

$$p = \frac{(\alpha\lambda - df)(1-q)}{2k}. \tag{5}$$

For ease of analysis and exposition, we rewrite the firm's expected payoff in (2),

$$\begin{aligned}\Pi_{attack} &= -[1-(1-t)(1-p)](1-q)\lambda - kq^2 \\ &= -[1-(1-t)(1-p)](1-q)\lambda - k(1-q)^2 + 2k(1-q) - k \\ &= -k(1-q)^2 + \{2k - [1-(1-t)(1-p)]\lambda\}(1-q) - k.\end{aligned} \tag{6}$$

Differentiating (6) with respect to $(1-q)$ and solving, we have

$$1 - q = \frac{2k - [1-(1-t)(1-p)]\lambda}{2k} = \frac{2k - t\lambda - (1-t)p\lambda}{2k} \tag{7}$$

Solving (5) and (7),

$$p^*_{attack} = \frac{(\alpha\lambda - df)(2k - t\lambda)}{4k^2 + (\alpha\lambda - df)(1-t)\lambda}, \tag{8}$$

and

$$1 - q^*_{attack} = \frac{2k(2k - t\lambda)}{4k^2 + (\alpha\lambda - df)(1-t)\lambda}. \tag{9}$$

Substituting in (6), the firm's maximum payoff without a bug bounty program is

$$\Pi^*_{attack} = \frac{4k^3(2k - t\lambda)^2}{[4k^2 + (\alpha\lambda - df)(1-t)\lambda]^2} - k. \tag{10}$$

Similarly, substituting (8) and (9) in (1), the strategic hacker's payoff is

$$U^*_{attack} = k\left[\frac{(\alpha\lambda - df)(2k - t\lambda)}{4k^2 + (\alpha\lambda - df)(1-t)\lambda}\right]^2. \tag{11}$$

4.2 Equilibrium With a Bug Bounty Program

Differentiating (3) with respect to p and solving, the strategic hacker's optimal choice of p when joining the bug bounty program is

$$p = \frac{R(1-q)}{2k}. \tag{12}$$

Rewrite the firm's payoff function in (4),

$$\begin{aligned}\Pi_{bounty} =& -\lambda t(1-p)(1-q) - kq^2 - p(1-q)R \\ =& -k(1-q)^2 + [2k - pR - \lambda t(1-p)](1-q) - k.\end{aligned} \tag{13}$$

Differentiating (13) with respect to $(1-q)$ and solving, we have

$$1 - q = \frac{2k - pR - t\lambda(1-p)}{2k}. \tag{14}$$

Solving (12) and (14),

$$p^*_{bounty} = \frac{(2k - t\lambda)R}{4k^2 + R(R - t\lambda)}, \tag{15}$$

and

$$1 - q^*_{bounty} = \frac{2k(2k - t\lambda)}{4k^2 + R(R - t\lambda)} \tag{16}$$

Substituting in (13), the firm's maximum payoff conditional on R is

$$\Pi_{bounty} = \frac{4k^3(2k - t\lambda)^2}{[4k^2 + R(R - t\lambda)]^2} - k. \tag{17}$$

By (17), the firm gets the highest payoff when $R = \frac{t\lambda}{2}$. However, for the strategic hacker to participate in the bug bounty program, his payoff must exceed the payoff obtainable by attacking the firm given the firm's choice of q. Substituting (5) in (1) and (12) in (3) and comparing, $U_{bounty} \geq U_{attack}$ if and only if $\frac{[R(1-q)]^2}{4k} \geq \frac{[(\alpha\lambda - df)(1-q)]^2}{4k}$, or

$$R \geq \alpha\lambda - df. \tag{18}$$

We need to consider the following two cases.

Scenario 1: $R = \frac{t\lambda}{2}$ satisfies (18), or $df \geq \lambda(\alpha - \frac{t}{2})$. We have $R^* = \frac{t\lambda}{2}$ and

$$\Pi^*_{bounty1} = \frac{4k^3(2k - t\lambda)^2}{\left[4k^2 - \frac{(t\lambda)^2}{4}\right]^2} - k. \tag{19}$$

Scenario 2: $R = \frac{t\lambda}{2}$ does not satisfy (18), or $df < \lambda(\alpha - \frac{t}{2})$. To incentivize the strategic hacker to join the bug bounty program, the firm must raise the reward. Therefore, we have a corner solution, $R^* = \alpha\lambda - df$. The firm's maximum payoff with the bug bounty program is then

$$\Pi^*_{bounty2} = \frac{4k^3(2k - t\lambda)^2}{[4k^2 + (\alpha\lambda - df)(\alpha\lambda - df - t\lambda)]^2} - k. \tag{20}$$

4.3 When Should a Bounty Program be Offered?

Summarizing the firm's payoff in the various scenarios:

- With no bug bounty program: $\Pi^*_{attack} = \frac{4k^3(2k-t\lambda)^2}{[4k^2+(\alpha\lambda-df)(1-t)\lambda]^2} - k$.
- With a bug bounty program and $df \geq \lambda(\alpha - \frac{t}{2})$: $\Pi^*_{bounty1} = \frac{4k^3(2k-t\lambda)^2}{\left[4k^2 - \frac{(t\lambda)^2}{4}\right]^2} - k$.
- With a bug bounty program and $df < \lambda(\alpha - \frac{t}{2})$: $\Pi^*_{bounty2} = \frac{4k^3(2k-t\lambda)^2}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^2} - k$.

The magnitudes of these payoffs depend on λ, d, f, α and t. In particular, $\Pi^*_{bounty2} < \Pi^*_{attack}$ when $\alpha > \frac{df}{\lambda} + 1$, meaning when α is sufficiently large, it is not optimal for the firm to offer a bug bounty program. Figure 1 depicts the optimal choices of the firm along different values of α and t. The strategic hacker prefers the bug bounty program when his expected benefit from joining the program exceeds that from launching a cyber attack, or $R \geq \alpha\lambda - df$.

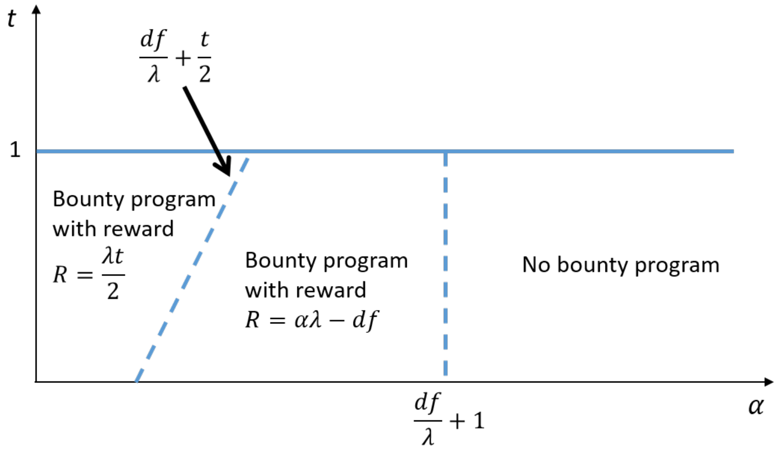
t
$\frac{df}{\lambda} + \frac{t}{2}$
1
Bounty program with reward $R = \frac{\lambda t}{2}$
Bounty program with reward $R = \alpha\lambda - df$
No bounty program
$\frac{df}{\lambda} + 1$
α

Figure 1: Optimal choice of bounty program

When α exceeds $\frac{df}{\lambda} + 1$, the firm is better off focusing on in-house protection. Recall the strategic hacker's opportunity cost of not attacking the firm is $\alpha\lambda - df$ (provided he has identified the vulnerability). To offset this opportunity cost, the bug bounty reward should increase with α to entice the strategic hacker. However, when α is too large, the cost to offer the reward will be excessive. The firm will prefer to shoulder the expected loss due to a vulnerability exploitation.

By contrast, when $\alpha \leq \frac{df}{\lambda} + 1$, the firm prefers to offer a bug bounty program. More importantly, when α is small and t is large, the firm will offer a relatively large bounty reward to incentivize the strategic hacker to compete with naive hackers in discovering the vulnerability. Our first proposition follows.

Proposition 1 *The firm will offer a bug bounty program if and only if $\alpha \leq \frac{df}{\lambda} + 1$. The bounty reward, $R = \frac{t\lambda}{2}$ when $df \geq \lambda(\alpha - \frac{t}{2})$, and $R = \alpha\lambda - df$ when $df < \lambda(\alpha - \frac{t}{2})$.*

In the remaining part of this paper, we assume $\alpha \leq \frac{df}{\lambda} + 1$ so that the firm optimally offers a bug bounty program. In addition, we mentioned in the model setting that k is sufficiently large to prevent perfect security[7]. Now we can more accurately define what "sufficiently large" means in this context. It means $q^*_{attack} < 1$, which implies $2k > t\lambda$.

Assumption 1 $\alpha \leq \frac{df}{\lambda} + 1$*, so that the firm optimally offers a bug bounty program.*

Assumption 2 $2k > t\lambda$*, so perfect security is not economical.*

4.4 Benefits of the Bug Bounty Program

To characterize the benefits brought by the bug bounty program, rewrite the firm's payoffs in (2) (with no bug bounty program) and (4) (with a bug bounty program):

$$\Pi_{attack} = (1-q)[-t(1-p)\lambda - p\lambda] - kq^2, \tag{21}$$

$$\Pi_{bounty} = (1-q)[-t(1-p)\lambda - pR] - kq^2. \tag{22}$$

In (21) and (22), conditional on the firm's effort, the first term in the square bracket is the expected loss when the vulnerability is identified by naive hackers. The second term is the expected loss or bug bounty reward payable when the vulnerability is identified by the strategic hacker. The third term is the firm's in-house effort cost.

By Proposition 1, when the firm prefers the bug bounty program, we must have $\alpha\lambda - df \leq \lambda$. Comparing the second term in the square bracket in (21) and (22), the firm can take advantage of the deterrence effect due to government enforcement, df, and the difference in her valuation of the system, λ, from that of the strategic hacker, $\alpha\lambda$, and offer a bug bounty reward, R, that is smaller than λ but yet enough to effectively dissuade the strategic hacker from attacking her (recall $R = \frac{t\lambda}{2} < \lambda$ when $df \geq \lambda(\alpha - \frac{t}{2})$ and $R = \alpha\lambda - df$ when $df < \lambda(\alpha - \frac{t}{2})$). We call this benefit of the bug bounty program **attack diversion**. It leads to a Pareto improved outcome as the firm does not need to suffer a full system loss and the strategic hacker can mostly receive a higher payoff because he does not need to face the potential legal penalty.

[7] similar for m as $m = k$ by assumption

The bug bounty program offers another important benefit. Comparing (8) with (15) and by Proposition 1, it is easy to see that $p^*_{bounty} \geq p^*_{attack}$. This means that with the bug bounty program, the strategic hacker will invest more effort to identify the vulnerability. Referring to the first term in the square bracket in (21) and (22), an increase in p will decrease the chance that naive hackers find the vulnerability. Hence, the bug bounty program can incentivize the strategic hacker to *compete* with naive hackers. It is akin to "recruiting" the strategic hacker to help "protect" the firm's system. We call this benefit of the bug bounty program **protection delegation**. By (21) and (22), the higher the threat due to naive hackers, t, the higher the benefit the firm can obtain from enlisting the strategic hacker's protection.

Furthermore, by (22) and Proposition 1, because the firm will offer $R < \lambda$, it prefers paying the bounty reward to having the system attacked by the strategic hacker. The firm can calibrate the strategic hacker's effort by modifying R. When the naive hacker's threat, t, is small, the firm's key concern is the attack from the strategic hacker. Hence, it will offer R just high enough to divert the strategic hacker's attack. As t increases, the threat due to naive hackers become prevalent. The firm will now try to induce more effort from the strategic hacker, p, by offering even higher bounty rewards so that the strategic hacker will compete with naive hackers. Figure 2 depicts how the bug bounty reward changes with the threat from naive hackers.

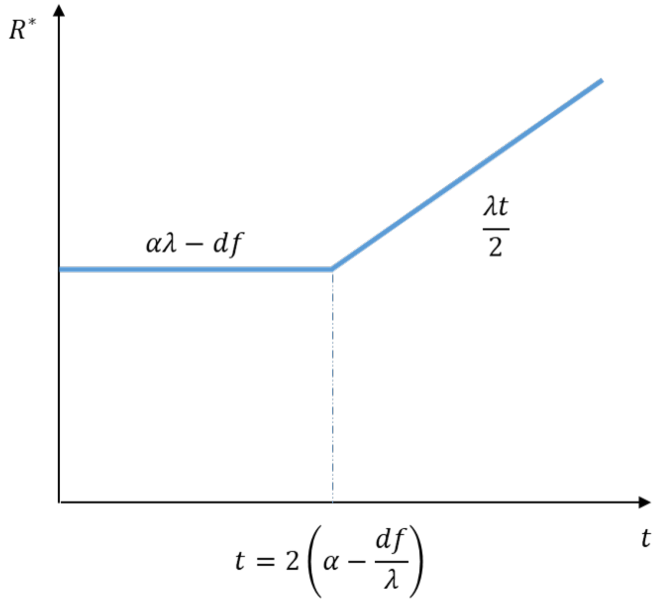

Figure 2: Optimal choice of bounty reward

Observation 1 *When the threat from naive hackers is small, i.e., $t < 2(\alpha - \frac{df}{\lambda})$, the firm will offer a constant bounty reward $R = \alpha\lambda - df$. The salient incentive is to divert attack away from the strategic hacker. When the threat from naive hackers is significant, i.e., $t \geq 2(\alpha - \frac{df}{\lambda})$, the firm will offer $R = \frac{\lambda t}{2}$. The salient incentive is to delegate the protection to the strategic hacker.*

4.5 Managing Bug Bounty Program

This section is concerned about how to integrate the bug bounty program as a component of comprehensive security strategies. Specifically, we seek to understand how the firm should balance between in-house protection and the bug bounty program. Should the firm, following the logic of IT outsourcing (e.g. Lacity et al. 2009), outsource the less-efficient in-house vulnerability identification activities (e.g. hiring security researchers, penetration test, security audits, etc.) to the "crowd", and focus on selecting and fixing the identified vulnerabilities?

Comparing (9) and (16), it is easy to see $q_{bounty} < q_{attack}$, meaning the bug bounty program does substitute in-house protection. However, by (16)

$$q_{bounty} = \frac{2kt\lambda - R(R - t\lambda)}{4k^2 + R(R - t\lambda)} = \begin{cases} \frac{2kt\lambda + (\alpha\lambda - df)(\alpha\lambda - df - t\lambda)}{4k^2 + R(R - t\lambda)} > 0, df < \alpha\lambda - \frac{\lambda t}{2} \\ \frac{2kt\lambda - \frac{(t\lambda)^2}{4}}{4k^2 + R(R - t\lambda)} > 0, df \geq \alpha\lambda - \frac{\lambda t}{2} \end{cases} \tag{23}$$

So $q_{bounty} > 0$ always hold. Therefore, while outsourcing some of the protection work to the "crowd"—the firm should never outsource all of them. Why would the firm (strategically) retain the in-house protection work, even though they do not have advantage in it by assumption?

To understand this, let us first examine the incentive for the firm to take in-house protection in a bug bounty program. By (4) we write the firm's best response function,

$$\frac{\partial \Pi_{bounty}}{\partial q} = \lambda t(1 - p) + pR - 2kq. \tag{24}$$

Therefore, the firm's optimal in-house protection effort depends on three factors, represented by three terms in the right hand side of (24): (1) the marginal loss due to the system attacked by naive hackers. The higher expected loss due to naive hackers, the greater incentive for the firm to strengthen in-house protection. (2) the marginal bounty reward payable to the strategic hacker. The higher the expected bounty payment, the greater the incentive for the firm to step up its protection effort (because it helps reduce the chance that it needs to pay the bounty reward.) (3) the marginal loss due to its own protection effort.

Having a more effective bug bounty program, i.e. higher p, affects the first two factors in an opposite way. On the one hand, it reduces the cost in the first term (i.e. the threat from naive hackers), so in terms of protecting the system against naive hackers, the necessity for the firm to undertake in-house protection does reduce. We call this strategic response **protection**

free-riding, as it arises when the firm wants to leverage the strategic hacker's effort in finding the vulnerability.

On the other hand, higher p increases the cost in the second term, hence the firm has incentives to increase its in-house protection to reduce the chance of paying the bounty reward. This is the key reason why firm should strategically retain in-house protection—it invests in self-protection in the hope that the cost will be justified by the reduced bounty payment. Intuitively, if a firm launches a bug bounty program with a very insecure system (e.g. zero in-house protection), the expected bug bounty payment could be excessive. A smarter strategy for the firm is to find and fix all the "easier" vulnerabilities. It is beneficially to do so as long as the cost of finding "easier" vulnerability can lead to even greater reduction in later bounty payments. We call this strategic response **bounty-payment squeezing**, as it arises when the firm wants to substitute bounty payment by in-house protection.

Proposition 2 *With a bug bounty program, the firm would retain in-house protection to squeeze the bug bounty payment.*

We next examine the strategic interaction (Bulow et al. 1985, Vives 2005) between the firm and the strategic hacker. The purpose is to explore how the bug bounty program changes the firm's and the strategic hacker's incentives to protect or attack the system.

4.5.1 No Bug Bounty Program

Differentiate (1) with respect to q and p, we have

$$\frac{\partial^2 U_{attack}}{\partial p \partial q} = -\frac{\partial^2 U_{attack}}{\partial p \partial (1-q)} = df - \alpha\lambda < 0, \tag{25}$$

which implies the strategic hacker's effort in identifying the vulnerability is a strategic substitute (Bulow et al. 1985) with the firm's effort. Intuitively, the higher effort the firm invests in in-house protection, the lower expected payoff the strategic hacker would expect from attacking the system, which causes him to scale down his investment.

Similarly, differentiate (2) with respect to p and q, we have

$$\frac{\partial^2 \Pi_{attack}}{\partial q \partial p} = -\frac{\partial^2 U_{attack}}{\partial (1-q) \partial p} = (1-t)\lambda > 0 \tag{26}$$

because $t < 1$. Hence, the firm's effort is a strategic complement (Bulow et al. 1985) with the strategic hacker's effort. Here, without a bug bounty program, the strategic hacker's gain is the firm's loss. Hence, when the strategic hacker increases his effort, the firm will respond by spending more effort to compete with the strategic hacker.

Proposition 3 *With no bug bounty program, the strategic hacker's effort is a strategic substitute with the firm's in-house protection. By contrast, the firm's in-house protection is a strategic complement with the strategic hacker's effort.*

4.5.2 With a Bug Bounty Program

Here, the strategic hacker's effort in identifying the vulnerability is affected by the firm's effort and the bounty reward. Differentiating (3),

$$\frac{\partial^2 U_{bounty}}{\partial p \partial q} = -\frac{\partial^2 U_{bounty}}{\partial p \partial (1-q)} = -R < 0, \tag{27}$$

$$\frac{\partial^2 U_{bounty}}{\partial p \partial R} = (1-q) > 0. \tag{28}$$

Hence, the strategic hacker's effort is again a strategic substitute with the firm's effort. However, his effort is a strategic complement with the bug bounty reward, R.

Differentiating the firm's payoff function in (4),

$$\frac{\partial^2 \Pi_{bounty}}{\partial q \partial p} = -\frac{\partial^2 U_{bounty}}{\partial (1-q) \partial p} = R - \lambda t. \tag{29}$$

Depending on the relative magnitude of R and λt, the firm's effort can be either a strategic substitute or a strategic complement with the strategic hacker's effort in identifying the vulnerability. From Figure 2, the only possibility for $R \geq \lambda t$ is when $t \leq 2(\alpha - \frac{df}{\lambda})$, i.e., the threat from naive hackers is small, in which case $R = \alpha\lambda - df$.

Taken together, for $R \geq \lambda t$ and $R = \alpha\lambda - df$, we must have $t \leq \alpha - \frac{df}{\lambda}$. Proposition 4 follows.

Proposition 4 *With a bug bounty program, the strategic hacker's effort in identifying the vulnerability is a strategic substitute with the firm's effort. When the threat from naive hackers is low, i.e., $t \leq \alpha - \frac{df}{\lambda}$, the firm's effort is a strategic complement with the strategic hacker's effort. When $t > \alpha - \frac{df}{\lambda}$, the firm's effort is a strategic substitute with the strategic hacker's effort.*

Figure 3 summarizes all the strategic interactions between the firm and the strategic hacker.

To understand the impact due to the naive hacker's threat, t, we refer back to (24). When the threat from naive hackers is significant, i.e., $t > \alpha - \frac{df}{\lambda}$, the first term tends to dominate the second term. It suggests when p increases, the gain from protection free-riding would be greater

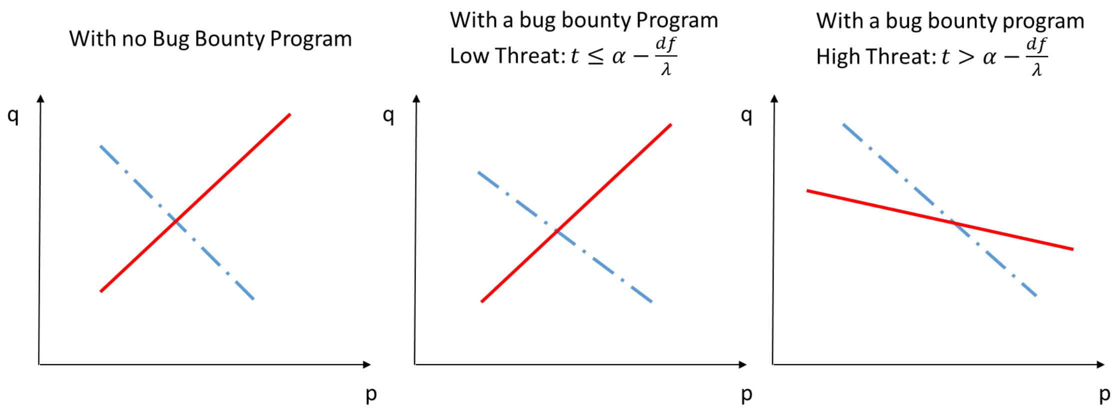

Figure 3: Strategic interactions (red solid line: best response of the firm, blue dotted line: best response of the strategic hacker)

than the loss due to the bounty payment, so the firm tends to reduce its own in-house effort. By contrast, when the threat from naive hackers is low, i.e., $t \leq \alpha - \frac{df}{\lambda}$, the bug bounty payment becomes more significant. The firm will prefer to increase its effort to squeeze the bounty payment in response to an increased p. Therefore, how the firm responds to the strategic hacker's effort depends on the external environment, viz. t.

4.6 Effect of Law Enforcement

We analyze how law enforcement affects the bug bounty program. For simplicity, we consider the expected net penalty due to law enforcement, df, together instead of separately considering the probability of apprehension, d, and the penalty term, f.[8]

Without a bug bounty program, by (8), (9), (10), and (11),

$$\frac{\partial p^*}{\partial df} < 0, \frac{\partial q^*}{\partial df} < 0, \frac{\partial U^*_{attack}}{\partial df} < 0, \frac{\partial \Pi^*_{attack}}{\partial df} > 0. \tag{30}$$

The interpretation is straightforward: law enforcement deters the strategic hacker by reducing his expected payoff. The firm can ride on this deterrence effect and reduce its in-house effort while enjoying a higher payoff. This result is consistent with classic economic theories which shows that public enforcement reduces crime (e.g., see Chalfin and McCrary 2017).

With a bug bounty program, by (16), (19), (20), and Proposition 1, the firm's payoff and

[8] In other settings, the probability of apprehension and the penalty term may produce different consequences in the analysis (e.g., Becker 1968).

its in-house effort are:

$$\Pi^*_{bounty} = \begin{cases} \frac{4k^3(2k-t\lambda)^2}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^2} - k, & df < \lambda(\alpha - \frac{t}{2}) \\ \frac{4k^3(2k-t\lambda)^2}{[4k^2-\frac{(t\lambda)^2}{4}]^2} - k, & df \geq \lambda(\alpha - \frac{t}{2}) \end{cases} \tag{31}$$

$$q^* = \begin{cases} 1 - \frac{2k(2k-t\lambda)}{4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)}, & df < \lambda(\alpha - \frac{t}{2}) \\ 1 - \frac{2k(2k-t\lambda)}{4k^2-\frac{(t\lambda)^2}{4}}, & df \geq \lambda(\alpha - \frac{t}{2}) \end{cases} \tag{32}$$

Differentiating with respect to df,

$$\frac{\partial \Pi^*_{bounty}}{\partial df} \begin{cases} \frac{8k^3(2k-t\lambda)[2(\alpha\lambda-df)-t\lambda]}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^3} > 0, & df < \lambda(\alpha - \frac{t}{2}) \\ = 0, & df \geq \lambda(\alpha - \frac{t}{2}) \end{cases} \tag{33}$$

$$\frac{\partial q^*}{\partial df} = \begin{cases} \frac{-2k(2k-t\lambda)[2(\alpha\lambda-df)-t\lambda]}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^2} < 0, & df < \lambda(\alpha - \frac{t}{2}) \\ = 0, & df \geq \lambda(\alpha - \frac{t}{2}) \end{cases} \tag{34}$$

(33) and (34) suggest that strengthening law enforcement increases the firm's payoff and decreases its in-house effort only when df is small. Recall from Observation 1 that when $df < \lambda(\alpha - \frac{t}{2})$, the optimal bounty reward is $\alpha\lambda - df$, which is determined based on the strategic hacker's incentive compatibility constraint, (18). In this case, the firm actually prefers a smaller bounty payment, $\frac{t\lambda}{2}$, but it has to offer a higher reward to entice the strategic hacker, which is to its disadvantage. Referring to Figure 4, the gap between $\alpha\lambda - df$ and $\frac{t\lambda}{2}$ is the value that the firm has to surrender to the strategic hacker. Strengthening the law enforcement, df, will decrease the strategic hacker's opportunity cost of submitting the vulnerability. This increases the "bargaining power" of the firm and helps decrease the bounty reward, which indirectly increases the firm's profit.

Furthermore, by the analysis following Proposition 4, for all $t \leq \alpha - \frac{df}{\lambda}$ which implies $df < \lambda(\alpha - \frac{t}{2})$, the firm has incentives to squeeze the bug bounty payment, R, by increasing its in-house effort, q. When the bounty reward is decreased due to law enforcement (see Figure 4), the firm can correspondingly cut back its in-house protection.

Note that law enforcement brings benefit to the firm only when $\frac{t\lambda}{2} \leq \alpha\lambda - df$. When $\frac{t\lambda}{2} > \alpha\lambda - df$, or $t > 2(\alpha - \frac{df}{\lambda})$ (see Figure 2), the firm's preferred bounty reward will already entice the strategic hacker to submit the vulnerability to the firm, so law enforcement does not matter any more. Therefore, $\frac{\partial \Pi^*_{bounty}}{\partial df} = 0$ and $\frac{\partial q^*}{\partial df} = 0$ when $df \geq \lambda(\alpha - \frac{t}{2})$. Similarly, law enforcement matters only when $\alpha\lambda - df \geq \frac{t\lambda}{2} \Rightarrow \alpha\lambda > \frac{t\lambda}{2}$, or $t < 2\alpha$. If $t \geq 2\alpha$, the firm will offer

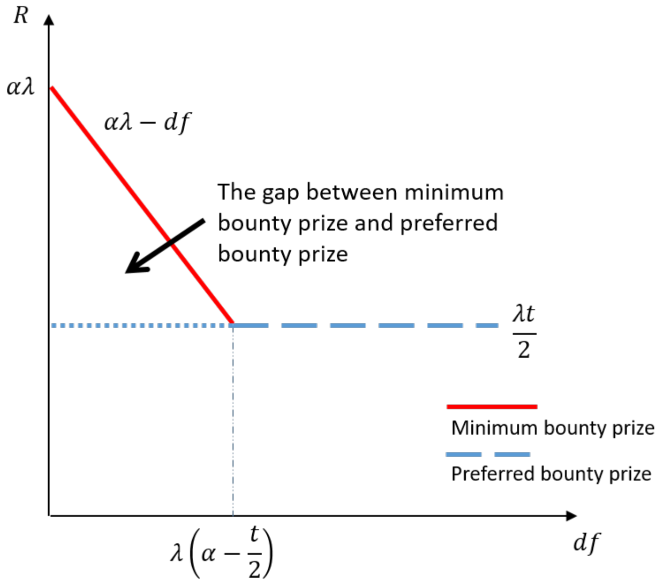

Figure 4: Law enforcement and bounty reward

a high bounty reward anyway, meaning law enforcement cannot help it further.

Next, by (3), (15), and Proposition 1,

$$U^*_{bounty} = \begin{cases} \frac{k(\alpha\lambda-df)^2(2k-t\lambda)^2}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-t\lambda-df)]^2}, & df < \lambda(\alpha-\frac{t}{2}) \\ \frac{k(t\lambda)^2(2k-t\lambda)^2}{4[4k^2-\frac{(t\lambda)^2}{4}]^2}, & df \geq \lambda(\alpha-\frac{t}{2}) \end{cases} \tag{35}$$

$$p^* = \begin{cases} \frac{(2k-t\lambda)(\alpha\lambda-df)}{4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)}, & df < \lambda(\alpha-\frac{t}{2}) \\ \frac{(2k-t\lambda)t\lambda}{2(4k^2-\frac{(t\lambda)^2}{4})}. & df \geq \lambda(\alpha-\frac{t}{2}) \end{cases} \tag{36}$$

Differentiating with respect to df,

$$\frac{\partial U^*_{bounty}}{\partial df} \begin{cases} (1-\frac{2k}{\alpha\lambda-df})\frac{2k(1+\frac{2k}{\alpha\lambda-df})(2k-t\lambda)^2(\alpha\lambda-df)^3}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^3} > 0, & 0 \leq df < max\{\alpha\lambda-2k,0\} \\ (1-\frac{2k}{\alpha\lambda-df})\frac{2k(1+\frac{2k}{\alpha\lambda-df})(2k-t\lambda)^2(\alpha\lambda-df)^3}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^3} \leq 0, & max\{\alpha\lambda-2k,0\} \leq df < \lambda(\alpha-\frac{t}{2}) \\ = 0, & df \geq \lambda(\alpha-\frac{t}{2}) \end{cases} \tag{37}$$

$$\frac{\partial p^*}{\partial df} = \begin{cases} (1-\frac{2k}{\alpha\lambda-df})\frac{(1+\frac{2k}{\alpha\lambda-df})(2k-t\lambda)(\alpha\lambda-df)^2}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^2} > 0, & 0 \leq df < max\{\alpha\lambda-2k,0\} \\ (1-\frac{2k}{\alpha\lambda-df})\frac{(1+\frac{2k}{\alpha\lambda-df})(2k-t\lambda)(\alpha\lambda-df)^2}{[4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)]^2} \leq 0. & max\{\alpha\lambda-2k,0\} \leq df < \lambda(\alpha-\frac{t}{2}) \\ = 0, & df \geq \lambda(\alpha-\frac{t}{2}) \end{cases} \tag{38}$$

When $df < \lambda(\alpha - \frac{t}{2})$, the analysis above suggests that strengthening law enforcement decreases the bounty reward and indirectly decreases the firm's in-house effort. These two effects produce opposite incentives for the strategic hacker. A smaller bounty reward decreases the strategic hacker's expected revenue in finding the vulnerability, so he responds by reducing his effort. However, a reduction in the firm's in-house effort increases the strategic hacker's expected gain, making him increase his effort. Depending on the balance between these two competing effects, law enforcement can variously increase or decrease the strategic hacker's equilibrium effort. Figure 5 depicts the changes in the optimal R, q, p, and U_{bounty} with df.

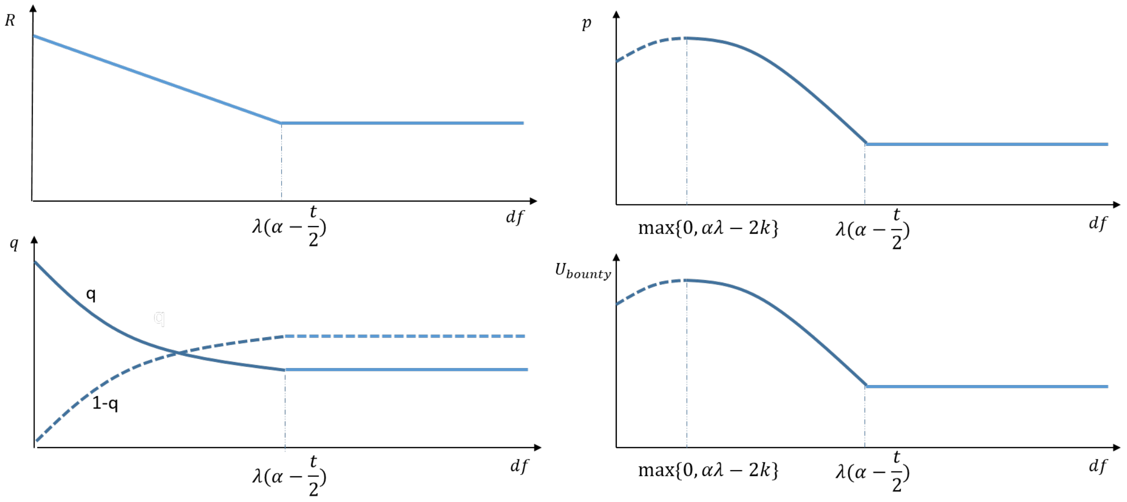

Figure 5: Effect of law enforcement

When $\alpha\lambda - 2k > 0$, the strategic hacker's payoff and effort increase first and then decrease with stronger law enforcement. In particular, when $df < \alpha\lambda - 2k$, the effect due to less in-house protection effect (lower q) dominates the effect due to reduced bounty reward (lower R), so the strategic hacker will increase his effort and get higher expected payoff. By contrast, when $\alpha\lambda - 2k \leq df < \lambda(\alpha - \frac{t}{2})$, the bounty reward effect dominates the in-house protection effect, so the strategic hacker will reduce his effort. When $df \geq \lambda(\alpha - \frac{t}{2})$, strengthening law enforcement does not affect the firm's strategy in the bug bounty program, and so it would not affect the strategic hacker either. Our next proposition follows.

Proposition 5 *Law enforcement does not have any effect on the firm, the strategic hacker, and the bug bounty program when $df \geq \lambda(\alpha - \frac{t}{2})$. When $df < \lambda(\alpha - \frac{t}{2})$, increasing law enforcement increases the firm's payoff and decreases its in-house protection effort. It increases (decreases) the strategic hacker's payoff and effort when $df < \alpha\lambda - 2k$ ($df \geq \alpha\lambda - 2k$).*

Proposition 5 provides an interesting insight —too much law enforcement, i.e., $df \geq \lambda(\alpha - \frac{t}{2})$, does not affect the strategic hacker because the bug bounty program offered by the firm can already dissuade the strategic hacker from attacking it. In fact, when the law enforcement lies in some middle range, i.e., $\alpha\lambda - 2k \leq df < \lambda(\alpha - \frac{t}{2})$, further increasing the enforcement could lead the firm and strategic hacker to spend *less* efforts in identifying the vulnerability. This increases the expected harm due to naive hackers and so makes the system *less secure*!

This counter-intuitive result arises because the enforcement has distorted the incentives of the firm and strategic hacker. Without the enforcement, the firm is obliged to pay a bigger bounty reward to the strategic hacker so that he will not attack the system. Strengthening the enforcement increases the "bargaining power" of the firm and so squeezes the bug bounty reward size. This causes both the firm and the strategic hacker to cut back their efforts. The firm invests less effort because its bug bounty cost is now reduced. The strategic hacker reduces its effort because he gets less reward from the bounty program.

Overall, the bug bounty reward, R, is an instrument for the firm to calibrate its own effort and the strategic hacker's effort in identifying the vulnerability. Enhancing law enforcement beyond $\alpha\lambda - 2k$ would disturb this calibration of R and lead to less secure protection of the system or produces no effect when it goes beyond $\lambda(\alpha - \frac{t}{2})$.

4.7 Are Bug Bounty Programs Economical?

Both Microsoft and Oracle criticized bug bounty programs based on cost-benefit arguments. For example, Oracle rejected bug bounty programs because of the belief that it is more economical to identify vulnerabilities themselves than to rely on external experts.

Our analysis so far partly responds to this criticism by showing that a bug bounty program can benefit a firm as long as cyber-attacks cost the firm more than what the hackers can gain from directly attacking the system. For this attack diversion benefit, whether the in-house protection is more cost-effective than the bug bounty program is irrelevant.

However, for the other type of benefit, viz. protection delegation, whether the in-house protection is more economical is not trivial. So far, we have assumed that the firm and the strategic hacker share exactly the same technology in identifying vulnerabilities. In this section, we relax this assumption and analyze how cost difference affects the firm's incentive to offer a bug bounty program. Such cost difference can come from, e.g., different expertise or experience in cyber-attack and protection.

Formally, by following the same analysis in the previous sections, we can derive the equilibrium strategies and payoffs for the firm and the strategic hacker when $m \neq k$.

With No bug bounty program

$$p^* = \frac{(\alpha\lambda - df)(2m - t\lambda)}{4mk + (\alpha\lambda - df)(1-t)\lambda}, \tag{39}$$

$$1 - q^* = \frac{2k(2m - t\lambda)}{4mk + (\alpha\lambda - df)(1-t)\lambda}, \tag{40}$$

$$\Pi^*_{attack} = \frac{4k^2m(2m - t\lambda)^2}{[4mk + (\alpha\lambda - df)(1-t)\lambda]^2} - m, \tag{41}$$

$$U^*_{attack} = \frac{k(\alpha\lambda - df)^2(2m - t\lambda)^2}{[4mk + (\alpha\lambda - df)(1-t)\lambda]^2}. \tag{42}$$

With a bug bounty program

$$p^* = \frac{R^*(2m - t\lambda)}{4mk + R^{*2} - t\lambda R^*}, \tag{43}$$

$$1 - q^* = \frac{2k(2m - t\lambda)}{4mk + R^{*2} - t\lambda R^*}, \tag{44}$$

$$\Pi^*_{bounty} = \frac{4k^2m(2m - t\lambda)^2}{(4mk + R^{*2} - t\lambda R^*)^2} - m, \tag{45}$$

$$U^*_{bounty} = \frac{kR^{*2}(2m - t\lambda)^2}{(4mk + R^{*2} - t\lambda R^*)^2}. \tag{46}$$

The company's net gain from the bug bounty program is

$$\Delta\Pi = \Pi_{bounty} - \Pi_{attack} = 4m(2m - t\lambda)^2 \left[\frac{1}{(4m + \frac{R^{*2} - t\lambda R^*}{k})^2} - \frac{1}{(4m + \frac{(\alpha\lambda - df)(1-t)\lambda}{k})^2} \right]. \tag{47}$$

By Observation 1, $R^* = max\{\frac{\lambda t}{2}, \alpha\lambda - df\}$. Hence, when $\alpha < \frac{df}{\lambda} + 1$, $4m + \frac{R^2 - t\lambda R}{k} < 4m + \frac{(\alpha\lambda - df)(1-t)\lambda}{k}$, meaning $\Delta\Pi > 0$. This is consistent with our earlier conclusion that whenever it is not too expensive to attract the strategic hacker, it is beneficial for the company to offer a bug bounty program. Meanwhile, the gain from the bug bounty program can be different depending on the cost coefficient of the strategic hacker, k. So, we need to analyze how k affects $\Delta\Pi$.

Differentiating $\Delta\Pi$ with respect to k,

$$\frac{\partial\Delta\Pi}{\partial k} = \frac{8m(2m - t\lambda)}{k^2} \left[\frac{R^{*2} - t\lambda R^*}{(4m + \frac{R^{*2} - t\lambda R^*}{k})^3} - \frac{(\alpha\lambda - df)(1-t)\lambda}{(4m + \frac{(\alpha\lambda - df)(1-t)\lambda}{k})^3} \right].$$

By proposition 1, there are two scenarios:

Scenario 1: $t \geq 2\left(\alpha - \frac{df}{\lambda}\right)$. We have $R^{*2} - t\lambda R^* = -\frac{(t\lambda)^2}{4}$, and so $\frac{\partial\Delta\Pi}{\partial k} < 0$. This means that the firm's gain from the bug bounty program is decreasing in k. When the exogenous threat due to naive hackers, t, is high, the firm always prefers smaller k, meaning more capable participants in the bug bounty program is good to the firm.

Scenario 2: $t < 2\left(\alpha - \frac{df}{\lambda}\right)$. We have $R^{*2} - t\lambda R^* = (\alpha\lambda - df)(\alpha\lambda - df - t\lambda)$. Let $k^* = \frac{(\alpha\lambda - df)(1-t)\lambda}{4m}\left[\frac{\alpha\lambda - df - t\lambda}{(1-t)\lambda}\right]^{\frac{1}{3}}\left[1 + \left(\frac{\alpha\lambda - df - t\lambda}{(1-t)\lambda}\right)^{\frac{1}{3}}\right]$, then it is easy to show that

$$k \leq k^* \Rightarrow \frac{\partial \Delta\Pi}{\partial k} > 0,$$

and

$$k > k^* \Rightarrow \frac{\partial \Delta\Pi}{\partial k} < 0.$$

This implies that the firm prefers a specific type of strategic hacker with cost k^*. Having more skilled participants than k^* actually decreases the firm's payoff. In particular, the negative relationship between k and m in k^* suggests that the firm prefers a complementary strategic hacker. When it has a cost advantage, i.e., m is small, it prefers a less capable participant, i.e., a participant with a large k, and vice versa.

When the firm has a cost advantage against the strategic hacker, the bug bounty program mainly helps it divert attacks away from the strategic hacker. The firm wants to spend less in the bug bounty program, and so it prefers less capable strategic hacker as it can then pay a lower expected bounty reward. By contrast, when the cost for the firm to identify the vulnerability is high, its main incentive is to delegate the protection to the strategic hacker. Hence, it prefers a more capable strategic hacker who can help it beat naive hackers in finding the vulnerability. When the exogenous threat due to naive hackers is salient, i.e., when $t \geq 2\left(\alpha - \frac{df}{\lambda}\right)$, the protection delegation incentive dominates the attack diversion incentive. Therefore, a more capable strategic hacker is better for the firm. Our next proposition follows.

Proposition 6 *The firm prefers a more capable strategic hacker when* $t \geq 2\left(\alpha - \frac{df}{\lambda}\right)$. *It prefers a complementary strategic hacker (i.e., a high cost firm prefers a low cost strategic hacker, and a low cost firm prefers a high cost strategic hacker) when* $t < 2\alpha - \frac{2df}{\lambda}$.

Proposition 6 suggests that the firm is less motivated to offer a bug bounty program when: (1) the external threat, t, is large but the external capability to identify vulnerabilities is weak, or (2) the external threat, t, is not large, but the external capability is not complementary to the internal capability. In both scenarios, the benefit of a bug bounty program may not be lucrative enough to motivate the firm to offer it.

Our analysis provides a sharp contrast between a bug bounty program and conventional crowd sourcing contests. In the crowd sourcing literature or, more generally, the literature related to contest theories (see, e.g., Terwiesch and Xu 2008, Fullerton and McAfee 1999, Konrad 2007), a contest organizer prefers more capable contributors and they would devise mechanisms to select more capable participants to join the contests. In the case of a bug bounty program, although more capable participants (strategic hackers) can help the firm protect its system better, it may cost the firm more in diverting the attacks from the participants themselves.

4.8 Security Protection

We next analyze how a bug bounty program affects the overall security of the system. We use S to denote the probability that the system is eventually exploited (Probability of exploited). Then, $1 - S$ represents the overall protection level of the system (Probability of protected). In the remaining part, we focus on analyzing S, then the conclusion about 1-S is straightforward.

With no bug bounty program, the system is exploited if either the strategic hacker or naive hackers find the vulnerability before the firm. Hence, $S_{attack} = [1-(1-t)(1-p_{attack})](1-q_{attack})$. Rewriting it, we have

$$S_{attack} = t(1 - p_{attack})(1 - q_{attack}) + p_{attack}(1 - q_{attack}) \tag{48}$$

The first term measures the probability of vulnerability exploitation by naive hackers. The second term measures the probability of vulnerability exploitation by the strategic hacker.

With a bug bounty program, the system will be exploited only by naive hackers. Hence,

$$S_{bounty} = t(1 - p_{bounty})(1 - q_{bounty}) \tag{49}$$

Comparing (48) and (49), it is straightforward to see that the bug bounty program changes the protection of the system in the following three ways:

1. Removing $p_{attack}(1-q_{attack})$, i.e., the bug bounty program diverts the strategic hacker away from launching an attack. This effect raises the overall protection of the system.
2. By (8), (15), and Proposition 1, $p_{bounty} > p_{attack}$, so the $t(1-p)$ factor is smaller in (49). Conditional on the firm's in-house protection, the competition between the strategic hacker and naive hackers is increased, which reduces the threat from naive hackers. This effect also raises the overall protection of the system.
3. By (9), (16), and Proposition 1, $q_{bounty} < q_{attack}$, which causes $t(1-p)(1-q)$ to increase. Hence, the bug bounty program causes the firm to cut back its in-house protection effort. This effect decreases the overall protection of the system.

Although counter-intuitive, the system could become *less* secure with the bug bounty programs provided that the third effect dominates the first two effects. Formally, we compare S_{bounty} and S_{attack}. With no bug bounty program,

$$\begin{aligned} S_{attack} &= [1-(1-t)(1-p_{attack})](1-q_{attack}) \\ &= [1-(1-t)(1-\frac{(\alpha\lambda-df)(2k-t\lambda)}{4k^2+(\alpha\lambda-df)(1-t)\lambda})]\frac{2k(2k-t\lambda)}{4k^2+(\alpha\lambda-df)(1-t)\lambda}. \end{aligned} \tag{50}$$

With a Bug Bounty Program,

$$\begin{aligned} S_{bounty} &= t(1-p_{bounty})(1-q_{bounty}) \\ &= t(1-\frac{(2k-t\lambda)R}{4k^2+R(R-t\lambda)})\frac{2k(2k-t\lambda)}{4k^2+R(R-t\lambda)} \\ &= \begin{cases} t(1-\frac{(2k-t\lambda)(\alpha\lambda-df)}{4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)})\frac{2k(2k-t\lambda)}{4k^2+(\alpha\lambda-df)(\alpha\lambda-df-t\lambda)}, df < \alpha\lambda-\frac{\lambda t}{2} \\ t(1-\frac{(2k-t\lambda)\frac{\lambda t}{2}}{4k^2-\frac{(\lambda t)^2}{4}})\frac{2k(2k-t\lambda)}{4k^2-\frac{(\lambda t)^2}{4}}, df \geq \alpha\lambda-\frac{\lambda t}{2}. \end{cases} \end{aligned} \tag{51}$$

The relative magnitude of S_{attack} and S_{bounty} depends on the parameters, meaning the bounty program need not provide better protection. Suppose that $k = 100$, $\lambda = 500$, $t = 0.2$, $\alpha = 0.5$. We have $t\lambda = 100$ and $\alpha\lambda = 250$. From (50) and (51):

$$S_{attack} = [1-0.8*(1-\frac{100(250-df)}{40000+400*(250-df)})]\frac{200*100}{40000+400*(250-df)} \tag{52}$$

and

$$\begin{aligned} S_{bounty} &= 0.2*(1-\frac{100*R}{40000+R(R-100)})\frac{200*100}{40000+R(R-100)} \\ &= \begin{cases} 0.2*(1-\frac{100*(250-df)}{40000+(250-df)(150-df)})\frac{200*100}{40000+(250-df)(150-df)}, df < 200 \\ 0.2*(1-\frac{100*50}{40000-2500})\frac{200*100}{40000-2500}, df > 200 \end{cases} \end{aligned} \tag{53}$$

Figure 6 plots the relative magnitudes of S_{attack}, S_{bounty} and $S_{bounty} - S_{attack}$ along different value of df.

In this numerical example, the bug bounty program leads to better overall protection only when df is extreme, i.e., when it is near 0 or $\alpha\lambda$. When law enforcement is moderate, having a bug bounty program is actually not good in terms of system protection!

The intuition is as follows. The difference between λ and $R = \alpha\lambda - df$ is the smallest when df is close to 0. By (9) and (16), q^*_{bounty} will be closest to q^*_{attack}, meaning we will have the least reduction in in-house protection. The first two security-enhancing effects characterized above due to the introduction of the bug bounty program will dominate the third security-decreasing effect,

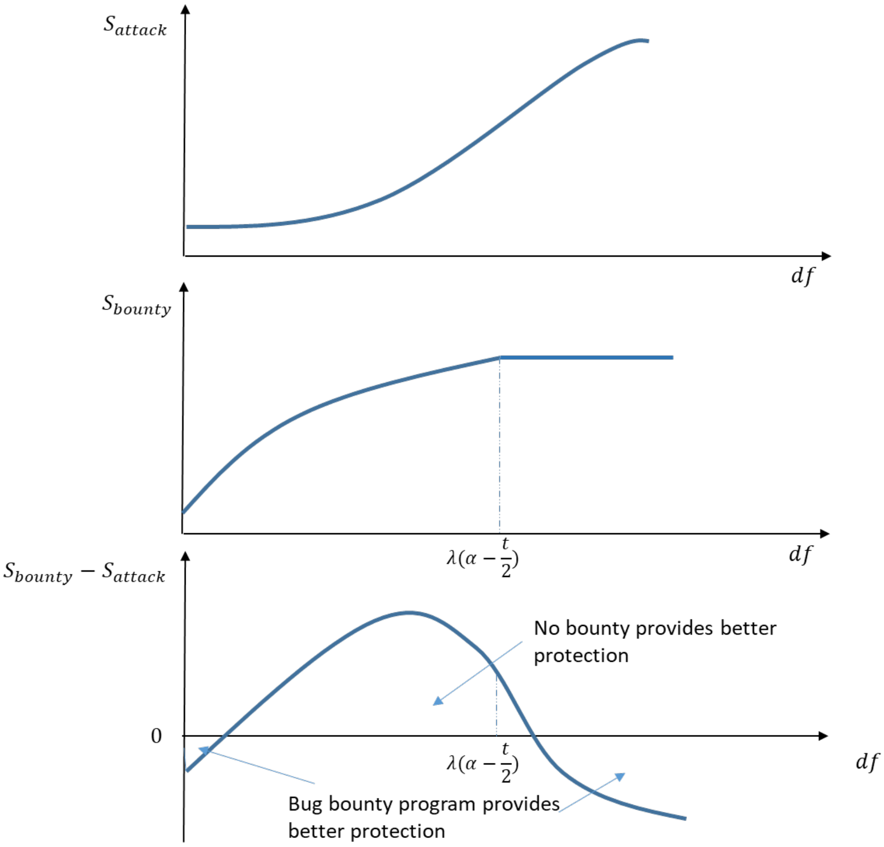

Figure 6: S_{attack}, S_{bounty}, $S_{bounty} - S_{attack}$ with df

leading to better overall protection.

As *df* increases, q_{bounty} gradually moves away from q_{attack}. Hence, the negative impact due to the reduction in in-house protection becomes more prominent until it dominates the first two security-enhancing effects characterized above. The bug bounty program then becomes undesirable from a system protection point of view.

When *df* becomes sufficiently large, in particular, when $df \geq \lambda(\alpha - \frac{t}{2})$, Proposition 5 suggests that p_{bounty} and q_{bounty} are no longer affected by law enforcement. Hence, the overall protection with a bug bounty program, $1 - S_{bounty}$, becomes constant. By contrast, without a bug bounty program, increasing *df* will continue to decrease q_{attack} which tends to worsen the protection. When this reduction in protection is large enough, the overall security level will become lower than that with a bug bounty program.

Observation 2 *A bug bounty program need not provide better protection.*

5 Conclusion

Despite its popularity, the merit of bug bounty programs is controversial, with many open questions ranging from how to design and operate bug bounty programs to whether they can fulfill the fundamental mission of security improvement. This study sheds light on these important issues. We develop a security game framework to analyze the trade-offs in offering a bug bounty program. We find that the bug bounty program is beneficial to a firm as long as it is not too costly for the firm to entice the strategic hacker to participate in the bounty program. With a bug bounty program, the firm enjoys two benefits: attack diversion and protection delegation. To get the best out of the bug bounty program, the firm should calibrate the bounty reward and in-house protection with reference to the severity of exogenous threat, reflecting the balance of two incentives, viz. bounty-payment squeezing and protection free-riding. Surprisingly, a firm should always retain in-house protection even though it is not more efficient in security protection, and bug bounty programs do not guarantee better security protection.

This study has important managerial implications. We show that the bug bounty program is not a a one-size-fits-all solution. Firms do need to evaluate their own security environment, the value and vulnerability of their systems, and in-house protection strategies to make better use of bug bounty programs. More importantly, our analysis clarifies some fundamental controversies about bug bounty programs, concerning its relationship with in-house protection and whether it actually brings economic and security benefits when the firms are variously competent/incompetent in security protection. We hope these clarifications can help firms make better decisions when considering bug bounty programs as a component of their comprehensive security strategies.

This work also provides some useful insight to law enforcement and public policies. As shown in our analysis, law enforcement interacts with private protection in a nuanced manner when a bug bounty program is introduced. We show that too much enforcement can reduce both in-house protection and the effectiveness of the bug bounty program (e.g., by reducing the probability that participants in the bug bounty program can identify the vulnerability), leading to even less secure system. As the bug bounty program is increasingly popular, the government can make use of this analysis in guiding its extent of intervention.

In terms of research, this study provides novel insights on security protection. Previous studies on security investment mostly focus on "sticks", i.e., how we could punish or deter hackers by lowering their expected pay from launching an attack (e.g., by strengthening in-house protection or punishment terms). Contrary to this literature, a bug bounty program acts more like a "carrot" and provides incentives to encourage hackers to hack ethnically. As demonstrated in our analysis, combining the “carrot” and “stick” is often superior to having the “stick” alone. It is promising for future research to focus more on “carrots” when deliberating about information security.

References

Bandyopadhyay T, Liu D, Mookerjee VS, Wilhite AW (2014) Dynamic competition in it security: A differential games approach. *Information Systems Frontiers* 16(4):643–661.

Becker GS (1968) Crime and punishment: An economic approach. *The economic dimensions of crime*, 13–68 (Springer).

Boudreau KJ, Lacetera N, Lakhani KR (2011) Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management science* 57(5):843–863.

Bulow JI, Geanakoplos JD, Klemperer PD (1985) Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political economy* 93(3):488–511.

Chalfin A, McCrary J (2017) Criminal deterrence: A review of the literature. *Journal of Economic Literature* 55(1):5–48.

Cremonini M, Nizovtsev D (2009) Risks and benefits of signaling information system characteristics to strategic attackers. *Journal of Management Information Systems* 26(3):241–274.

Egelman S, Herley C, Van Oorschot PC (2013) Markets for zero-day exploits: Ethics and implications. *Proceedings of the 2013 New Security Paradigms Workshop*, 41–46 (ACM).

Fbogov (2018) Crowdsourced vulnerability discovery and disclosure (cvdd) services. https://www.fbo.gov/index?s=opportunity&mode=form&tab=core&id=7a142ab942e23617c0005dc9b2a717a6, retrieved Jan 29.

Finifter M, Akhawe D, Wagner D (2013) An empirical study of vulnerability rewards programs. *USENIX Security Symposium*, 273–288.

Fisher D (2010) Microsoft says no to paying bug bounties. https://threatpost.com/microsoft-says-no-paying-bug-bounties-072210/74249/, retrieved Jan 29.

Fryer H, Simperl E (2017) Web science challenges in researching bug bounties. *Proceedings of the 2017 ACM on Web Science Conference*, 273–277 (ACM).

Fullerton RL, McAfee RP (1999) Auctionin entry into tournaments. *Journal of Political Economy* 107(3):573–605.

Gal-Or E, Ghose A (2005) The economic incentives for sharing security information. *Information Systems Research* 16(2):186–208.

Gao X, Zhong W, Mei S (2015) Security investment and information sharing under an alternative security breach probability function. *Information Systems Frontiers* 17(2):423–438.

Gordon LA, Loeb MP (2002) The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* 5(4):438–457.

Hausken K (2006) Income, interdependence, and substitution effects affecting incentives for security investment. *Journal of Accounting and Public Policy* 25(6):629–665.

Horton JJ, Chilton LB (2010) The labor economics of paid crowdsourcing. *Proceedings of the 11th ACM conference on Electronic commerce*, 209–218 (ACM).

Jeppesen LB, Lakhani KR (2010) Marginality and problem-solving effectiveness in broadcast search. *Organization science* 21(5):1016–1033.

Kannan K, Telang R (2005) Market for software vulnerabilities? think again. *Management science* 51(5):726–740.

Keller J (2018) Vulnerability reward program: 2017 year in review. https://security.googleblog.com/2018/02/vulnerability-reward-program-2017-year.html, retrieved Jan 29.

Konrad KA (2007) Strategy in contests-an introduction .

Lacity MC, Khan SA, Willcocks LP (2009) A review of the it outsourcing literature: Insights for practice. *The Journal of Strategic Information Systems* 18(3):130–146.

Maillart T, Zhao M, Grossklags J, Chuang J (2017) Given enough eyeballs, all bugs are shallow? revisiting eric raymond with bug bounty programs. *Journal of Cybersecurity* 3(2):81–90.

Munaiah N, Meneely A (2016) Vulnerability severity scoring and bounties: Why the disconnect? *Proceedings of the 2nd International Workshop on Software Analytics*, 8–14 (ACM).

Oraclecom (2015) No you really cannot. https://gist.github.com/michaeldyrynda/9b5eac6c02e6089052a6, archived version. Retrieved Jan 29.

Png IP, Wang QH (2009) Information security: Facilitating user precautions vis-à-vis enforcement against attackers. *Journal of Management Information Systems* 26(2):97–121.

Ruohonen J, Allodi L (2018) A bug bounty perspective on the disclosure of web vulnerabilities. *Workshop on Economics of Information Security.*

Szymanski S (2003) The economic design of sporting contests. *Journal of economic literature* 41(4):1137–1187.

Terwiesch C, Xu Y (2008) Innovation contests, open innovation, and multiagent problem solving. *Management science* 54(9):1529–1543.

Vives X (2005) Complementarities and games: New developments. *Journal of Economic Literature* 43(2):437–479.

Zhao M, Grossklags J, Liu P (2015) An empirical study of web vulnerability discovery ecosystems. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1105–1117 (ACM).

Zhao M, Laszka A, Grossklags J (2017) Devising effective policies for bug-bounty platforms and security vulnerability discovery. *Journal of Information Policy* 7:372–418.