Harnessing Uncertainty in Vulnerability Market

Zhen Li
Department of Economics and Management
Albion College, USA
Email: zli@albion.edu

Qi Liao
Department of Computer Science
Central Michigan University, USA
Email: liao1q@cmich.edu

***Abstract*—Zero-day vulnerabilities pose significant threats in computer and network security, and have attracted attentions in recent years not only to malicious attackers but government and law enforcement users who need to control (e.g., for forensics purpose) the computer systems which otherwise are inaccessible through traditional channels. Based on the observation that vulnerabilities are acquired and traded in a different way than commodities, we study and propose a vulnerability market model by taking into consideration cheating and uncertainty in the market. The paper illustrates the interactions between the vulnerability sellers and buyers in a game theoretic framework. By modeling the economic aspects of the vulnerability market with a focus on information asymmetry and distinctive incentives of malicious and defensive buyers, we propose active and strategic market participation by defenders to obtain vulnerability information from the marketplace in a cost-effective way. Rather than killing the market, defenders can take advantage of the incomplete information feature of the vulnerability market to improve cyber-security. To further maximize the uncertainty, defenders may also play in the supply side of the vulnerability market to provide low or no value vulnerabilities to dilute the market.**

***Index Terms*—Computer Security, Zero-day Vulnerability Market, Uncertainty, Asymmetric Information, Game Theory, Economics**

I. Introduction

A zero-day vulnerability is a computer program flaw unknown to program vendor that may expose the program to malicious attack. Such security hole may then be exploited by hackers to take control of the systems for various purposes. Uses of zero-day exploits include infiltrating malware, spyware, ransomware or allowing unwanted access to user information. The value of zero-days comes from the fact that they are undisclosed to the software vendors and users. Whoever with the knowledge of a zero-day vulnerability can exploit it until the vendor or the user learns of it and has the hole patched. It is found that a typical zero-day attack lasts for 312 days on average before details of the holes surface in public, i.e., the average vulnerability window of a zero-day exploit is about 10 months [1].

There are various ways for vulnerability finders to disclose zero-day vulnerabilities: to disclose it publicly, to disclose it to the vendor, or to other interested parties. In recent years, the most common way of disposing vulnerability information is to sell it in a marketplace [2]. The market value of a new vulnerability ranges between $5,000-$250,000, assuming an exclusive sale, the most modern version of the software, and, not alerting the vendor. Some fees might even be paid in installments, with each subsequent payment depending on the vendor not patching the vulnerability used by the exploit [3]. Notably, the zero-day vulnerability market consists of three categories: the white market, in which vulnerabilities are sold to software vendors or other companies that help the developers rectify security flaws; the black market, where zero-day exploits are sold to criminal organizations and malicious users; and the intermediate gray market, where exploits are sold to legitimate buyers such as governments and vendors of espionage and monitoring Trojans [4]. Currently, vendors largely depend on their own bug bounty programs [5] to obtain vulnerability information (white-market transactions). Vendors and other defensive buyers intending to fix the vulnerability flaw are lack of active market participation.

We argue that unlike other commodities traded in conventional markets, vulnerability markets have their unique characteristics and are fundamentally different from traditional sense of markets. Being a market trading merely information, the vulnerability market is prevalent with hidden and asymmetric information. It is hard to assess the reliability and value of a vulnerability, to price a vulnerability, to secure the exclusive right of the buyer to the vulnerability, or to prevent the buyer from reselling the vulnerability information. Cheating is easy and common for both the seller and the buyer.

In the context of information asymmetry prevailing the vulnerability market, we ask the question "how can defenders take advantage of such information defect in the vulnerability market to improve cybersecurity?" To address the question, we study the economic aspects of the marketplace for zero-day vulnerabilities in a game theoretic framework. We analyze the strategic interactions between sellers and buyers of vulnerabilities, and explore the effects of cheating and uncertainty in the marketplace.

Based on the game theoretic analysis, we provide some thoughts on the issue of utilizing the information asymmetry on the vulnerability market to improve cybersecurity. First, malicious buyers and defensive buyers have different desire on exclusive access to vulnerability. To maximize the exploit benefits from the vulnerability, an attacker desires for exclusive knowledge of the vulnerability, and pays a premium to disincentive the seller from selling to multiple sources. Compared to attackers, vendors and other defensive buyers intending to fix the security holes have no similar desire for excludability. Thus they possess a cost advantage to attackers to purchase vulnerabilities from the marketplace. Currently, the market for vulnerabilities is under-utilized by defensive buyers as a way to obtain vulnerability information. Our findings suggest more participation of defenders playing a buying role in vulnerability markets for better cybersecurity.

Second, given the uncertain nature of the true quality of vulnerability information, we further propose defenders' participation playing a selling role in vulnerability markets. Defensive buyers may offer low-value or even fake vulnerabilities to the marketplace to confuse malicious buyers and dilute the market price for high-value vulnerabilities. Defenders may offer test trials of made-up vulnerabilities with the help of virtualization environment enabled by techniques such as honeypots. By introducing additional uncertainty, defenders have a better chance of obtaining high-quality vulnerabilities at the cost of malicious buyers. Either way, when defensive buyers join the demand side or the supply side of the vulnerability market, the odds against malicious buyers of obtaining high-value vulnerability from the marketplace increase.

The rest of the paper is organized as follows. Section II reviews related works. Section III identifies the key characteristics of the vulnerability market and emphasizes the uncertainties due to asymmetric information. In Section IV, we build the game theoretic framework for analyzing the strategic interactions between the seller (finder) of a vulnerability and malicious buyers (attackers) in the benchmark model where defenders are absent. We discuss the model implications of the players' best responses when the game is nonrepetitive or repeated. The welfare implications of pricing and selling strategies are also discussed. We extend the benchmark model in Section V to study how defensive buyers (defenders) may join both the demand side and the supply side of the vulnerability market to decrease vulnerability leakage to malicious buyers by taking advantage of asymmetric information prevailing the market. Section VI provides further discussions and implications of the modeling analysis. Section VII concludes the paper.

II. Related Work

In general, there has been limited literature on the zero-day vulnerability markets largely due to the difficulty of identifying and studying such markets [6]. Nevertheless, the discovery, dissemination and disclosure of vulnerabilities have been studied to some extent. For example, the effects of information disclosure are studied from the perspective of analysis of markets for sharing security information [7]. Although an early mathematical study of the vulnerability market proposed a federally-funded mechanism for software vulnerability disclosure [8], the real world vulnerability markets turn out to be much more complicated and unregulated. A majority of vulnerability finders are found not affiliated with the software vendors so that multiple vulnerability markets have emerged. Mostly the vulnerability market participants remain anonymous. Although there have been efforts to analyze the identities, motivation, and behaviors of vulnerability finders and buyers, the dataset is of limited size, and representative information is still missing [5].

Vulnerability markets are closely related to online illicit markets, many of which are used to trade stolen information, hacking tools, and other illegitimate resources. Disruption and intervention mechanisms have been studied aiming at disenabling such illicit online markets [9]. Regulatory approaches have been explored that may be effective in preventing acts associated with illegal online market transactions, including regulations relating to payment mechanisms used by market participants to make purchases and receive payments. Plausible payment regulations may be used to reduce payment providers' involvement in illegal online transactions [10].

To further counter illicit online markets, idea of market for lemons has been introduced. Since the distribution of relevant information, especially the information of true quality of the products, is asymmetric between the seller and the buyer, the buyer can be ripped by the seller due to lack of information. Trust-based disruption techniques such as Sybil and slander attacks have been proposed to create quality uncertainty, e.g., by leaving fictitious feedbacks [11]. Considering sellers with peaches may attempt to implement quality signaling to single out their quality products, law enforcement may engage in actual sales coupled with tracing buyers for arrest [12]. Either way, the market demand shrinks.

While the above regulatory methods or lemonizing techniques could also be applied to the vulnerability market to disincentive market participants, it is not as clear-cut so far whether it is optimal to eliminate the online market for vulnerability information. We explore the alternative solution that vulnerability markets may be desirable from the perspective of improve cybersecurity. There are both legal and illicit zero-day vulnerability trade, and legitimate vulnerability markets have been created to compete with black markets [13]. Empirical examinations of market-based vulnerability disclosure mechanism found that markets for vulnerabilities may be effective at restricting the diffusion of vulnerability exploitation, the risk of exploitation, and the volume of exploitation attempts [14].

There has been a growing debate over the role of the government in the zero-day market. Some researchers have suggested that the federal government corner the market, purchasing all known zero-days and revealing them [15] while others want to regulate the market and make the sale of zero-days illegal [4]. Our early work proposes the creation of incentive mechanisms for vendors to invest in security and encourages usage of the vulnerability market for defensive buyers to improve cybersecurity of governments and smart cities [16]. Attempts to either monopolize or restrict the zero-day market to specific parties are, nevertheless, likely not only to fail but also to undermine security by hindering legitimate research. Given the possible benefits gained from markets for vulnerabilities, rather than attempting to further lemonize the vulnerability market, we study the plausible mechanisms to take advantage of the asymmetric information prevailing the vulnerability market to reduce vulnerability information leakage to bad actors. We argue that the key is to increase the likelihood such information for sale in the marketplace falls in the hands of defensive, rather than offensive buyers.

TABLE I: Characterization of Goods

	Rival	Non-Rival
Excludable	Private Goods	Natural Monopolies
Non-Excludable	Common Resources	Public Goods

III. Challenges in the Vulnerability Marketplace

In this section, we describe the key characteristics of the zero-day market, and discuss the asymmetric information problems prevailing the market.

A. *Vulnerability Information Not As Commodity*

In economics, goods are commonly classified according to two features: excludability and rivalry. A good is excludable if consumption can be detected and prevented by the seller. A good is rival if one person's use of it diminishes other people's use. Depending on whether a good is excludable or rival, it is placed in one of the four categories, as in Table I.

A market can be formed for any good that is excludable, but only private goods can be efficiently produced and allocated in the marketplace. Most information goods such as software and weather reports are non-rival and can be shared by large population. They are considered public goods when freely accessible, and natural monopolies when they can only be accessed with fee.

Vulnerability information traded in the zero-day market, nevertheless, does not fit in any common category. It is excludable (one has to buy it in order to use it) and rival (a buyer's benefit is reduced when sharing the vulnerability information with other buyers), but it is not a private good. Different from a physical good, the transfer of the vulnerability information from the seller to a buyer does not secure the exclusive use of the information by the buyer. It is difficult, if not impossible, to prevent the seller from selling to multiple buyers. Actually, the buyer, after acquiring the information, may also choose to resell the information to one or multiplier buyers. Therefore, vulnerability information traded in the marketplace does not belong to any established category of defined commodities. As the vulnerability sold to multiple parties can be considered as a common resource, the zero-days are some mixture of private goods and common resources.

B. *Uncertainties in the Zero-day Market*

Vulnerability exploiters have two ways to acquire vulnerability information, to explore for vulnerabilities on their own or to purchase vulnerability information from the marketplace. Self exploration can be time consuming, and exploiters largely depend on the vulnerability market for information. Vulnerability discoverers and exploiters are two separate groups [5].

Due to the nature of vulnerability information and zero-day exploits, as specified in Section III-A, there are many obstacles in the vulnerability market that trades in traditional goods and services do not have to face. Most importantly, the vulnerability market is full of uncertainties.

First, the value of vulnerability is unclear and unstable. Factors that can affect its value include the exploit ability of the vulnerability, the timing of the patch, how many parties know about the vulnerability, etc. Vulnerabilities are not interchangeable nor directly comparable thus pricing is difficult. Without an effectively-functioned pricing mechanism, it is hard for market participants to value the vulnerability information to price it appropriately. One side will always lose out.

Second, the information of vulnerability is hidden to buyers. The details of the vulnerability have to be kept secret before selling. Often it is difficult even to accurately describe a vulnerability without making the vulnerability easier to find. There is lack of transparency in terms of the same information accessible to all market participants. The zero-day market is subject to extortion where buyers can be ripped off by sellers.

Last but not least, the exclusive rights of the vulnerability information cannot be guaranteed. In order to receive the largest payoff from exploitation, the exploiter must be willing to have the exclusive access to all rights to the vulnerability. However, the vulnerability seller may offer the same information to numerous sources, and buyers have no effective ways to prevent multiple selling or even to tell whether multiple selling has occurred. In fact, sellers adopt a business model that often plays their customers against one another as they try to keep up in an espionage arms race [17]. Agencies would have been eager to pay more for exclusive use of the vulnerability, but there is no way to ensure that sellers will not sell to multiple buyers. There is also no way to ensure that some buyers will not reveal or resell their exploits. The value of secrecy complicates the efforts to control the vulnerability trade because it contributes to market opacity and lack of transparency about buyer and seller behavior.

Vulnerability markets can be labeled as a lemon market, in which there is quality uncertainty therefore those selling quality products are unable to differentiate themselves from sellers with poor quality products, and cannot compete with their low prices. Engaging in a lemon market increases the effort and cost of buyers and reduces their expected benefits. Participating in the vulnerability market is even more challenging for buyers because they not only face the uncertainty in product quality, but also the uncertainty pertaining whether the sale is exclusive or inclusive.

IV. The Benchmark Vulnerability Game Model

In this section, we build the game theoretic framework to study the interactions between sellers and malicious buyers of vulnerabilities in a benchmark model where defensive buyers are not present. Best strategies of the two sides of the market are discussed when the game is either nonrepetitive or repetitive. One particular uncertainty in the vulnerability market is modeled, i.e., unknown reselling information.

A. *Game Players and One-Shot Payoffs*

We use a game theoretic framework to analyze the interactions between a vulnerability discoverer and a number of potential buyers who are interested in obtaining the vulnerability. Vulnerability discoverers include hackers and security researchers who find and sell vulnerabilities. In the benchmark model, all potential buyers are malicious. They do not report the vulnerability information purchased, but they intend to use the vulnerabilities to compromise systems. Zero-days are especially valuable to them when kept unknown to software vendors, and vulnerabilities are most exploitable when kept secret to other interested parties as well, in particular competing exploiters.

Suppose there exist N potential buyers competing with one another, who buys the vulnerability for immediate or future exploit use. The value of the vulnerability to each buyer depends on the likelihood for the exploit to succeed and the payoff of successful exploit of buyer's own assessment, both are decreasing in the number of exploiters possessing the information. In particular, as the vulnerability information is known to more exploiters, the chance for the software vendor to be aware of the vulnerability increases, thus decreasing the success rate of exploiting the vulnerability; the more attackers exploit the vulnerability, the less payoff is to receive for each exploiter. Therefore, regardless of the intrinsic value of the vulnerability, the vulnerability information is the most worthwhile to a buyer if the buyer has exclusive rights to it. A buyer's willingness to pay decreases if the buyer has to share the vulnerability information with multiple buyers.

When the seller negotiates the price with each buyer individually, the price Buyer i is willing to pay is $p_i(n)$, where $n \in [1, N]$, which is the actual number of exploiters sharing the vulnerability. $p_i(n)$ is decreasing in n. The seller needs to decide to how many buyers to sell the vulnerability to maximize revenue:

$$\max_{n \in [1,N]} \sum_{i=1}^{n} p_i(n)$$

Let $b_i(n)$ be the benefit received by Buyer i from exploiting the vulnerability when sharing the vulnerability with $(n-1)$ other buyers. b_i is decreasing in n. The buyer's goal is to maximize consumer surplus from acquiring vulnerability information from the marketplace:

$$\max_{p_i(n)} b_i(n) - p_i(n)$$

When $n = 1$, the user has the exclusive rights to the vulnerability information, and the benefit received with exclusive use is $b_i(1)$. Assume for any arbitrary buyer i, $b_i(1) - p_i(1) > b_i(n) - p_i(n)$, $\forall n \in [2, N]$, where the difference on each side of the inequality measures the consumer surplus received by Buyer i with or without the exclusive rights to the vulnerability. Therefore, an attacker, as a potential buyer of some vulnerability information, always prefers to having exclusive access to the information.

Note in the modeling setup, the seller determines how many buyers have possession of the vulnerability, but the price of the vulnerability depends on buyers' willingness to pay, which in turn depends on to what level the vulnerability is shared.

Proposition 1. In a complete information market, the vulnerability information is sold exclusively to Buyer 1 if $b_1(1) \geq \sum_{i=1}^{n} p_i(n)$ for any $2 \leq n \leq N$; the information is sold to multiple buyers if otherwise.

In a market or game with complete information, the payoffs, strategies and types of players are common knowledge. For our vulnerability-selling game in particular, the seller would be honest and credible. Buyers would face no risk of cheating. We rank the buyers based on the payoff expected to receive from exploit from the highest to the lowest (i.e., $b_1(n) > b_2(n) > ... > b_N(n)$) and call the buyer with the highest expected benefit Buyer 1. Since a buyer is willing to pay for the vulnerability if and only if the benefit of exploit exceeds the cost of acquiring the vulnerability, the rank of potential buyers' willingness to pay is the same as the rank of expected benefits. When the seller is honest when releasing relevant information, in particular to how many buyers the vulnerability information is sold to, the range of mutually accepted price for the seller to sell to Buyer 1 exclusively instead of to n buyers is $\sum_{i=1}^{n} p_i(n) \leq p_1(1) \leq b_1(1)$. Similarly, the range of mutually accepted price for the seller to sell to Buyer 2 exclusively is $\sum_{i=1}^{n} p_i(n) \leq p_2(1) \leq b_2(1)$, and so on. Since Buyer 1 expects to receive the largest benefit from exclusive access of the vulnerability information, Buyer 1 must be the actual buyer of the information in case of exclusive sale.

The seller's choice is whether to sell the vulnerability to Buyer 1 at a price bound by $b_1(1)$ or n buyers at a price capped by $\sum_{i=1}^{n} p_i(n)$. Thus, the sale must be exclusive if the former is bigger, and multiple if the latter is bigger. This is the optimal solution as it maximizes combined gains from trade between the seller and all potential buyers.

Proposition 2. In an incomplete information market, the seller's best strategy is to sell the vulnerability information to all potential buyers when the game is nonrepetitive.

When the multiple selling information is unknown to buyers, the optimal vulnerability selling, as specified in Proposition 1, cannot be realized. Although buyers can request for exclusive use of the vulnerability information, there is no security measures to prevent the seller from cheating. If the seller cheated and sold the vulnerability to multiple buyers ($n \geq 2$) at exclusive prices, then the revenue received would be $\sum_{i=1}^{n} p_i(1) > p_1(1)$, which is more financially attractive than playing honest. Clearly, the potential gain to the seller is the maximum when the vulnerability information is sold to all potential buyers while claiming to each buyer the sale is exclusive.

B. Solutions to the Benchmark Model

Based on the previous section, if we consider the game between the seller and a particular buyer, the short-term (one round of the game) payoff matrix is as shown in Table II. The buyer's strategy space is to demand exclusive access to the information or agree to share the information with $n-1$ other buyers at the time of purchase. The seller's strategy is to be honest or dishonest with the buyer regarding the number of buyers obtaining the vulnerability. If the seller chooses

TABLE II: Payoff Matrix of the Game between the Seller and One Buyer, Known Only to Seller under Incomplete Information

Buyer i Seller	Exclusive	Sharing
Honest	$p_i(1)$, $b_i(1) - p_i(1)$	$\sum_{i=1}^{n} p_i(n)$, $b_i(n) - p_i(n)$
Cheat	$\sum_{i=1}^{N} p_i(1)$, $b_i(N) - p_i(1)$	$\sum_{i=1}^{N} p_i(n)$, $b_i(N) - p_i(n)$

to cheat, the vulnerability information must be sold to all potential buyers to generate the largest possible revenue to the seller. In particular, cheating means that the seller secretively sells the same vulnerability to as many buyers as possible without being aware by any particular buyer. Note that when cheating is possible, the buyer is always worse off.

The buyer cannot tell whether the seller cheats or not. Given model assumptions, having exclusive rights to the vulnerability information generates the highest expected payoff to the buyer so that the buyer would have to choose to trust the seller at the beginning of the game (either the game is one-stage or in the first round of a repeated game). The seller has complete information by contrast and the full payoff table, as in Table II, is known only to the seller.

The payoff for the seller is higher when the seller plays the cheating strategy regardless how the buyer chooses. Hence, cheating dominates playing honest in the game. The seller will always choose to cheat if the game is non-repetitive. Hence in the one-stage benchmark game with incomplete information, when buyers have to trust the seller, the solution of the strategy combination would be $(cheat, exclusive)$, i.e., the seller will cheat by selling the same vulnerability information to all N buyers, and each buyer purchases vulnerability by asking for exclusive accessibility. Apparently, the seller dominates the game because of the information advantage.

C. Repetitive Game

The game between sellers and buyers in the vulnerability market can be infinitely repeated, when sellers are treated as a group of actors continuously supplying vulnerabilities to the marketplace. Selling a new vulnerability is considered as starting a new round of the game. The market starts with trust. As all buyers prefer to having exclusive access to a vulnerability, they will choose to trust the seller at the beginning round of the game and pay the high price for exclusive access.

If only one vulnerability is ever sold in the zero-day market, the seller must choose to cheat and sell to all potential buyers while promising them exclusive use and grabs the highest revenue possible ($\sum_{i=1}^{N} p_i(1)$). As vulnerability discoverers continue to supply vulnerabilities to the market, the game between sellers and buyers become repeated. If sellers are found cheat, buyers may retaliate in the future by cutting their willingness to pay.

Suppose after the N buyers have each purchased the vulnerability information in the first round of the game and started using it, each buyer has the probability of $\theta(n)$ of finding the existence of $n-1$ other co-users, and $\sum_{n=1}^{N} \theta(n) = 1$. In the next round of the game, i.e., when another vulnerability is provided to the market, each buyer will expect the same seller behavior as in the previous round of the game, the so-called "adaptive expectation." Therefore, in the second round of the game, Buyer i's willingness to pay is $\sum_{n=1}^{N} \theta_i(n) \times p_i(n)$. The expected payment by all N buyers received by the vulnerability seller in round two of the game is $\sum_{i=1}^{N}(\sum_{n=1}^{N} \theta_i(n) \times p_i(n))$, so is the expected payoff of the seller in all the following rounds.

Since today's payoff of \$1 is more valuable than tomorrow's \$1, we define $\delta \in (0, 1)$ as the discount factor. When the seller cheats in every round of the game, the total expected payoff of the seller in this infinitely repeated game is

$$\sum_{i=1}^{N} p_i(1) + (\delta + \delta^2 + \delta^3 + ...)\{\sum_{i=1}^{N}(\sum_{n=1}^{N} \theta_i(n) \times p_i(n))\}$$

Or equivalently,

$$\sum_{i=1}^{N} p_i(1) + \frac{\delta}{1-\delta}\{\sum_{i=1}^{N}(\sum_{n=1}^{N} \theta_i(n) \times p_i(n))\} \tag{1}$$

If the seller is always honest instead, the vulnerability will always be sold to Buyer 1 that values vulnerability information the highest, and the seller's total expected payoff is

$$(1 + \delta + \delta^2 + ...) \times p_1(1)$$

Or equivalently,

$$\frac{1}{1-\delta} \times p_1(1) \tag{2}$$

Proposition 3. An increasing premium is necessary for buyers to induce exclusive sale of a vulnerability.

To secure permanent exclusive access to all the vulnerability information, the lump-sum threshold premium paid by Buyer 1 must be

$$\sum_{i=2}^{N} p_i(1) + \frac{\delta}{1-\delta}\{\sum_{i=1}^{N}(\sum_{n=1}^{N} \theta_i(n) \times p_i(n)) - p_1(1)\} \tag{3}$$

The premium works to "bribe" the seller to behave. Obviously, as the demand pool increases, that is, as N increases, the premium required for exclusive sale also increases, making it hard to secure the exclusive access to the vulnerability information.

D. Pricing and Market Efficiency

The seller's goal is to maximize the revenue from selling vulnerability. Individual buyer's goal is to maximize the consumer surplus from buying the vulnerability. For the buyer, this means to secure the exclusive access to the vulnerability information. Exclusive access can only be guaranteed (i.e., the seller will be kept financially from cheating) when the buyer is willing to pay a premium no less than the threshold

TABLE III: Distribution of Vulnerability Information

Scenario	Defender	Attacker
I	don't know	don't know
II	know	don't know
III	don't know	know
IV	know	know

specified by Equation 3. This is only the case if the single buyer's willingness to pay exceeds the probabilistic payoffs of all other buyers combined.

Extending from one seller to numerous sellers, in any period, the normalized price of a vulnerability is

$$(1-\delta)(1+\delta+\delta^2+...)\{\sum_{i=1}^{N}(\sum_{n=1}^{N}\theta_i(n)\times p_i(n))\}$$

Or equivalently,

$$\sum_{i=1}^{N}(\sum_{n=1}^{N}\theta_i(n)\times p_i(n)) \tag{4}$$

This can be approximately considered as the normalized market price. It is associated with two major uncertainties: the uncertainty in the selling status of the vulnerability, and the uncertainty in buyers' willingness to pay. The latter is arguably related to the intrinsic value of the vulnerability information to each buyer.

Proposition 4. With incomplete information, market efficiency cannot be achieved in the vulnerability information market.

Market efficiency is achieved with the sum of consumer surplus and producer surplus, called "total surplus", being maximized. The calculation of total surplus is

$$\sum_i(b_i(n)-p_i(n))+(\sum_i p_i(n)-C)=\sum_i b_i(n)-C \tag{5}$$

where C is the seller's cost associated with finding and selling the vulnerability. Note in Equation 5, the price paid by the buyers and received by the sellers are canceled out.

Given the sellers' cost, total surplus depends on the benefits received by all buyers combined from exploiting the vulnerability. Nevertheless, the uncertainties in the zero-day market make it impossible to find the optimal size of n. That is, it is impossible to reach market efficiency in the real world with incomplete information.

Market efficiency is not equivalent to social optimum. In the market analysis, only the well-being of sellers and buyers of the vulnerability is considered. For the zero-day market, when buyers use the purchased information to attack, users of the affected software will be harmed.

V. Defenders' Active Participation in Vulnerability Markets

In this section, we extend the benchmark model to a three-party game with the involvement of the defenders in the vulnerability market. We study how defenders may take advantage of the lack of information in the vulnerability market to improve cybersecurity.

A. Uncertainties in Three-party Game

In the benchmark model, we assume all potential buyers are malicious with the purpose of acquiring vulnerability to attack. To prevent vulnerability from falling in the wrong hands, we allow the defenders to join the vulnerability market based on the principle that defenders may participate in the marketplace to compete with exploiters for information.

The uncertainties prevailing the zero-day vulnerability market are extended as follows:

- The market is anonymous. Identities of market participants are undisclosed.
- The quality of vulnerability information is unknown to buyers. Test trials and quality signaling are imperfect if existing.
- To how many buyers the seller is releasing vulnerability information is unknown to buyers.

There are four possible scenarios regarding who has access to the vulnerability information, as shown in Table III. In Scenario I, if there is a potential vulnerability on a target system, both the defender and the attacker do not know such vulnerability. In Scenario II, the defender knows the system has a vulnerability but the attacker does not know. It is the opposite in Senario III, where the attacker knows a specific vulnerability on the target system but the defender does not know that. Lastly, Scenario IV specifies the moment when both the defender and the attacker know the disclosed vulnerability. Apparently, Scenario III is what the attacker attempts to achieve, and is the worst case for the defender, who must prevent it from occurring.

To prevent Scenario III from happening, the defender may join the demand side and the supply side of the vulnerability market to manipulate information to serve the defense purpose. Figure 1 illustrates the circular flow of vulnerabilities and payments in the marketplace among market participants at the presence of defenders joining either the demand or the supply side (or both sides) of the vulnerability market, as discussed in the following sections.

B. Defender Joining the Demand Side of the Vulnerability Market

The vulnerability finder seeking to monetize the finding may share it with responsible disclosure programs and get the reward, sell on the black market but facing potential criminal prosecution, or arrange a deal through an exploit broker. Either way the finder is putting the vulnerability for sale in the marketplace. Ultimate buyers of the vulnerability information can be defensive or offensive. Unlike offensive buyers who intend to use the vulnerability to exploit, defensive buyers intend to defend the product against cyber attack. They will use the purchased vulnerability information to patch the product and make it more secure.

Vendors are already participating in the marketplace as defensive buyers, largely through vulnerability reward programs, but the participation is limited. Currently, there are only a few vulnerability reward programs by vendors, most of which were

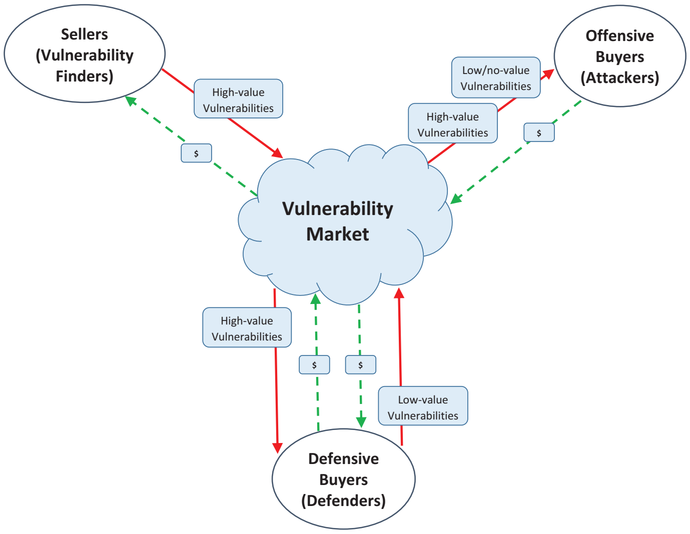
Sellers (Vulnerability Finders)
Offensive Buyers (Attackers)
High-value Vulnerabilities
Low/no-value Vulnerabilities
High-value Vulnerabilities
$
$
Vulnerability Market
High-value Vulnerabilities
$
$
Low-value Vulnerabilities
Defensive Buyers (Defenders)

Fig. 1: The model of vulnerability market with participation of defensive buyers, offensive buyers and sellers. The participants may have multiple roles (e.g., defensive buyers may switch to defensive sellers).

created a few years ago [5]. There are several reasons why selling vulnerabilities to vendors can be attractive, including the decreased risk of getting ripped off and the possibility of future job offers. Finders also receive recognition. The vulnerability reward programs were found to be economically efficient, comparing favorably to the cost of hiring full-time security researchers to locate bugs internally [18].

To improve cybersecurity, we propose active market participation by defensive buyers at the demand side to reduce the chance vulnerability information is purchased by offensive buyers. As Table III illustrates, the defender's goal is to avoid Scenario III. Therefore, it is not necessary for the defender to have exclusive access to the vulnerability information. All that matters is for the defender to be aware of the existence of the security holes so actions can be taken to patch the holes. Different from an attacker who may have to pay a premium to request exclusive sale, facing the risk of being cheated, the defensive buyer does not desire exclusive sale.

In the context of the benchmark model, we make the following specifications:

- The defender and the attacker both participate in the demand side of the vulnerability market.
- The defender and the attacker have different willingness to pay, depending on whether to request exclusive ownership.
- The seller either cannot tell the true characteristics of a buyer, being defensive or offensive, or do not care about the nature of a buyer.
- The seller sells to an exclusive buyer or multiple buyers to maximize revenue.

The defender's strategy space as a buyer falls in the range of $[p_d(N), p_d(1)]$, where the subscript d stands for the defensive buyer, and N is the number of potential buyers competing in the market, including the defensive buyer. The actual price paid by the defender would be the lowest possible price to acquire the vulnerability information. In the one-stage game, the defensive buyer purchases the vulnerability information at the lowest possible price $p_d(N)$. When the game is repeated, the presence of defensive buyers increases competition on the demand side of the vulnerability market which would increase the premium an exploiter has to pay to induce exclusive sale by the seller. As sellers have inherent financial incentives to sell to multiple sources, defensive buyers have the cost advantages than offensive buyers in the vulnerability market to purchase vulnerabilities.

C. Defender Joining the Supply Side of the Vulnerability Market

It is interesting to note that defenders may also join the supply side of the vulnerability market to dilute the vulnerability products supplied to the market. Defenders may intentionally supply low-value vulnerabilities to confuse exploiters. Furthermore, defenders may even supply fake vulnerabilities when quality signaling is imperfect.

With defenders joining the supply side of the vulnerability market, there are two types of sellers in the market: vulnerability finders (offering high-value vulnerabilities) and defenders (offering no-value or low-value vulnerabilities). Due to the uncertainties in vulnerability markets, defensive sellers may copy quality signaling offered by vulnerability finders to add more uncertainties to the marketplace. With deployment of dedicated honeypot farm and sophisticated virtual machine technologies, the boundary between real and fake (or artificial) vulnerabilities becomes to blur. The various uncertainties prevent malicious buyers from distinguishing between finders and defenders as sellers or between vulnerabilities.

Let b^h be the average benefit received from exploiting a high-value vulnerability, b^l be the average benefit received from exploiting a low-value vulnerability, and no benefit result from a fake vulnerability. Suppose of all the vulnerability products supplied to the market, ρ^h fraction is of high value and ρ^l fraction is of low value. Of the $(1-\rho^h-\rho^l)$ fraction of no-value vulnerabilities, let s fraction be claimed as high-value and $(1-s)$ fraction be claimed as low-value.

The defenders' goal is to minimize the probability for the offensive buyers to obtain high-value vulnerabilities:

$$\min_s \frac{\rho^h}{\rho^h + s \times (1-\rho^h-\rho^l)} \tag{6}$$

The best strategy is to provide both low-value vulnerabilities (ρ^l) and fake vulnerabilities disguised as high-value vulnerabilities. At $s = 1$, the probability for a malicious buyer to acquire a true vulnerability of high-value is

$$\frac{\rho^h}{1-\rho^l} \tag{7}$$

Since $\rho^h + \rho^l < 1$, $\frac{\rho^h}{1-\rho^l} < 1$. Malicious buyers exploiting high-value vulnerabilities are worse off at the presence of fake vulnerabilities introduced by defenders.

At the presence of low-value vulnerabilities provided by defenders, malicious buyers buy both high-value and low-value vulnerabilities. Treating all malicious buyers as an entity, the expected benefit received by a representative malicious buyer when purchasing a vulnerability from the market is

$$Eb = \rho^h \times b^h + \rho^l \times b^l \quad (8)$$

In the absence of defenders participating in the supply side, a representative malicious buyer receives the benefit of b^h. Since $\rho^h + \rho^l < 1$ and $b^l < b^h$, $Eb < b^h$. Malicious buyers are worse off for sure. The degree of welfare deterioration for attackers depends on the level of market dilution generated by defenders with low-value and no-value vulnerability products.

As a direct outcome of worse-off situation, malicious buyers' desire to participate in the market shrinks. For malicious buyers remaining in the market, part of their resources flow to defenders for purchasing low-value and no-value vulnerabilities, thus reducing the financial gains of vulnerability finders compared to the benchmark model.

The increased uncertainties are only applicable to offensive buyers and not to the defensive buyers. By playing sellers' role to introduce additional uncertainties to the vulnerability market, defenders increase the opportunity for themselves to purchase high-quality vulnerabilities, thus to reduce the chance for attackers to acquire high-value vulnerabilities, given attackers' limited information and financial constraints.

Lastly, besides creating financial disincentives to attackers, providing low-value vulnerabilities allows defenders to better defend. On one hand, defenders can concentrate their resources and be well prepared to defend against attacks exploiting specific low-value vulnerabilities leaked by defenders themselves. On the other hand, such mechanism also makes it possible for defenders to detect attackers' activities, an idea similar to honey-passwords [19].

VI. Discussion and Future Work

It is arguable whether the zero-day vulnerability market is beneficial to the security community [20]. Depending on the result of the debates, various methods may be designed according to our game theoretic analysis to either improve or to deteriorate the zero-day market. For example, monitoring mechanisms may be created to support the market. On the other hand, defenders may inject fake vulnerability information to the market to increase market frictions for malicious buyers, and hinder the smooth functioning of the vulnerability market. We posit that we should acknowledge the existence of the vulnerability market and explore how defenders may take advantage of the asymmetric information feature of the vulnerability market to obtain vulnerabilities from the marketplace.

Our analytic results suggest high correlation between pricing/premium and uncertainties in the vulnerability market. First, the market price of a vulnerability depends not only on the intrinsic value of the vulnerability, but also the information uncertainty associated with the vulnerability. Second, the premium to secure the exclusive use of the vulnerability information can be significantly high, and the distortion in pricing lies in the uncertainties in the zero-day market. Our model predicts a significantly high premium the buyer has to pay for exclusive access to the vulnerability information. The high premium the buyer is willing to or is forced to pay functions as a financial incentive to prevent the seller from cheating, but it is a double-edged sword. It can also induce the seller to cheat if the credibility of the seller cannot be testified. The effectiveness of defenders' participation in the vulnerability market largely depends on the cheating tendency of vulnerability finders.

Various uncertainties prevailing the vulnerability market are considered in this study, and the resulting inefficiency of the zero-day market is discussed as well. There is no easy solution to improve efficiency of the vulnerability market due to the unique characteristics of the information product sold in the market, which does not fit readily in any category of goods and services defined in economics. According to the modeling analysis, various factors matter to the uncertainty level of the zero-day market. Defenders may take measures to manipulate the factors to increase the uncertainty level of the market to depress the market.

Our analysis is limited to the interactions between the vulnerability information seller and the potential buyers of the seller, offensive and defensive. A broad understanding of the market may result from pursuit of a precise, mathematical model of more interdependent actors including customers of the software vendor, the government as buyer and regulator, etc. Our future work includes empirical research assessing the application of the proposals in marketplaces such as studies on exclusability, vulnerability pricing, and effect of fake vulnerability on the market. The further research will help improve our knowledge of the functioning of the vulnerability market, the impact of the proposed methods, and inspire us of future research on the zero-day vulnerabilities.

VII. Conclusion

Zero-day vulnerabilities have been a serious threat to modern day computer security. Parties ranging from malicious users, software vendors, government agencies, etc., are striving to acquire high-value vulnerability information in order to gain advantages from compromising target systems, whether for financial benefits or cyberspace warefares. Although vulnerabilities have been traded in the past, the market for zero-day vulnerabilities has not been studied as thoroughly as other aspects of the underground economy.

We take the effort to analyze the strategic interactions between sellers and buyers of vulnerabilities, and study the effects of cheating and uncertainty in the market. In particular, we study the determination of the price of vulnerability information in a game theoretic framework involving vulnerability finders (sellers), attackers (offensive buyers) and defenders (defensive buyers and sellers). We showed the winning buyer has to pay a significant premium to secure the exclusive access to the information. There are no effective measures to prevent the seller from cheating, but incentives may be created to affect the likelihood for the seller to cheat.

We believe the vulnerability market is currently underutilized by the defenders as a way to obtain vulnerability information, and propose active market participation by defenders, in both the demand side and the supply side, to reduce the risk of vulnerability exploitation. Defenders can increase the uncertainties of the zero-day market, resulting in a significant decrease of the likelihood for malicious buyers to acquire vulnerabilities. Different preferences regarding the exclusive access to vulnerability information between malicious buyers and defensive buyers give defenders financial advantages to obtain the information at relatively low cost. Supplying low-quality (or even fake) vulnerabilities to the market further helps defenders to dilute the market and increases the chance for defenders to obtain high-quality vulnerabilities. The proposed market participation methods take advantage of the vulnerability market imperfections to achieve cybersecurity goals.

References

[1] L. Bilge and T. Dumitras, "Before we knew it: an empirical study of zero-day attacks in the real world," in *Proceedings of the ACM Conference on Computer and Communications Security*, Raleigh, NC, October 16-18 2012, pp. 833–844.

[2] C. Miller, "The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales," in *Proceedings of the Sixth Workshop on the Economics of Information Security*, Pittsburgh, PA, June 2007.

[3] A. Greenberg, "Shopping for zero-days: A price list for hackers' secret software exploits," *Forbes*, March 23 2012.

[4] P. N. Stockton and M. Golabek-Goldman, "Curbing the market for cyber weapons," *Yale Law & Policy Review*, vol. 32, no. 1, pp. 101–128, 2013.

[5] A. M. Algarni and Y. K. Malaiya, "Software vulnerability markets: Discoverers and buyers," *International Journal of Computer, Information Science and Engineering*, vol. 8, no. 3, pp. 71–81, 2014.

[6] L. Allodi, "Economic factors of vulnerability trade and exploitation," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, Dallas, Texas, October 30-November 03 2017, pp. 1483–1499.

[7] A. Ozment, "Bug auctions: Vulnerability markets reconsidered," in *Proceedings of the Third Workshop on the Economics of Information Security*, Minneapolis, MN, May 13-14 2004.

[8] K. Kannan and R. Telang, "Market for software vulnerabilities? Think again," *Management Science*, vol. 51, no. 5, pp. 726–740, 2005.

[9] T. J. Holt, "Identifying gaps in the research literature on illicit markets on-line," *Global Crime*, vol. 18, no. 1, pp. 1–10, 2017.

[10] A. Hutchings and T. J. Holt, "The online stolen data market: disruption and intervention approaches," *Global Crime*, vol. 18, no. 1, pp. 11–30, 2017.

[11] J. Franklin, V. Paxson, S. Savage, and A. Perrig, "An inquiry into the nature and causes of the wealth of internet miscreants," in *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, Virginia, October 29-November 02 2007, pp. 375–388.

[12] S. C. Hoe, M. Kantarcioglu, and A. Bensoussan, "A game theoretical analysis of lemonizing cybercriminal black markets," in *Proceedings of the 3rd International Conference on Decision and Game Theory for Security*, Budapest, Hungary, November 5-6 2012, pp. 60–77.

[13] M. Zhao, J. Grossklags, and P. Liu, "An empirical study of web vulnerability discovery ecosystems," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, Denver, Colorado, October 12-16 2015, pp. 1105–1117.

[14] S. Ransbotham, S. Mitra, and J. Ramsey, "Are markets for vulnerabilities effective?" *Management Information Systems Quarterly*, vol. 36, no. 1, pp. 43–64, 2012.

[15] A. Schwartz and R. Knake, "Governments role in vulnerability disclosure: Creating a permanent and accountable vulnerability equities process," *The Cyber Security Project*, 2016.

[16] Z. Li and Q. Liao, "Economic solutions to improve cybersecurity of governments and smart cities via vulnerability markets," *Government Information Quarterly*, vol. 35, no. 1, pp. 151–160, January 2018.

[17] A. Greenberg, "Meet the hackers who sell spies the tools to crack your pc (and get paid six-figure fees)," *Forbes*, March 21 2012.

[18] M. Finifter, D. Akhawe, and D. Wagner, "An empirical study of vulnerability reward programs," in *Proceedings of the 22nd USENIX Conference on Security*, Washington, D.C., August 14-16 2013, pp. 273–288.

[19] A. Juels and R. L. Rivest, "Honeywords: making password-cracking detectable," in *Proceedings of the ACM SIGSAC conference on Computer & communications security (CCS)*, Berlin, Germany, November 4-8 2013, pp. 145–160.

[20] S. Egelman, C. Herley, and P. C. van Oorschot, "Markets for zero-day exploits: ethics and implications," in *Proceedings of the New Security Paradigms Workshop*, Banff, Alberta, Canada, September 09-12 2013, pp. 41–46.