MIS Quarterly

RESEARCH ARTICLE

ARE MARKETS FOR VULNERABILITIES EFFECTIVE?[1]

Sam Ransbotham
Carroll School of Management, Boston College, Chestnut Hill, MA 02467 U.S.A. {sam.ransbotham@bc.edu}

Sabyaschi Mitra
College of Management, Georgia Institute of Technology, Atlanta, GA 30332 U.S.A. {saby.mitra@mgt.gatech.edu}

Jon Ramsey
SecureWorks, Inc., P.O. Box 95007, Atlanta, GA 30347 U.S.A. {jramsey@secureworks.com}

Current reward structures in security vulnerability disclosure may be skewed toward benefitting nefarious usage of vulnerability information rather than responsible disclosure. Recently suggested market-based mechanisms offer incentives to responsible security researchers for discovering and reporting vulnerabilities. However, concerns exist that any benefits gained through increased incentives for responsible discovery may be lost through information leakage. Using perspectives drawn from the diffusion of innovations literature, we examine the effectiveness of market-based vulnerability disclosure mechanisms. Empirical examination of two years of security alert data finds that market-based disclosure restricts the diffusion of vulnerability exploitations, reduces the risk of exploitation, and decreases the volume of exploitation attempts.

Keywords: Information security, vulnerability disclosure, information technology policy

Introduction

Historically, relatively simple, disconnected systems with little data provided few security concerns. In recent years, dramatic changes in information systems complexity, connectedness, and economic value have led to greatly increased security concerns (Baskerville 1993; Dhillon and Backhouse 2001; Lohmeyer et al. 2002; Sandhu and Samarati 1996; Straub and Welke 1998). Consequently, the equity markets penalize the announcement of software vulnerabilities (Cavusoglu et al. 2004; Telang and Wattal 2007), and recent legislation such as Sarbanes-Oxley provides penalties for inadequate security (Schultz 2004). Clearly, information security is important in the current business environment (Gordon and Loeb 2002).

Attackers often compromise computer systems by exploiting vulnerabilities present in the software running on these systems (Cavusoglu et al. 2007; Cavusoglu et al. 2008). The impact of a software vulnerability depends on who makes the discovery. If discovered first by the technology vendor or benign security professionals, they have the opportunity to eliminate the vulnerability or otherwise protect systems before they are attacked; if discovered first by potential attackers, or if vulnerability patches and countermeasures are not installed quickly after disclosure, then systems are at risk. From this perspective, vulnerability discovery is a race between attackers and security professionals where "incentives are becoming as important as technical design" in protecting systems (Anderson and Moore 2006, p. 610). Unfortunately, incentives in this environment may be skewed to benefit attackers. While attackers are compensated through using (or selling) vulnerabilities (Radianti and Gonzalez 2007), security researchers have not historically been paid for discovering vulnerabilities.

[1]H. Raghav Rao wass the accepting senior editor for this paper. Wei T. Yue served as the associate editor.

In recent years, mechanisms such as vulnerability markets (Kannan and Telang 2005; Schechter 2004) and auctions (Ozment 2004) have been proposed to provide incentives to security professionals to discover vulnerabilities. Several private vulnerability markets currently exist, including the *Vulnerability Contributor Program* from *iDefense* (Verisign), and the *Zero Day Initiative* from *Tipping Point Technologies* (3COM). These markets pay an undisclosed price for each verified vulnerability reported to them based on its potential impact (Kannan and Telang 2005). They also incorporate the vulnerability in the intrusion detection systems (IDS) and other countermeasures that they provide to their subscribers for immediate protection, notify the vendor about the vulnerability so that the vendor can develop and propagate patches to correct the vulnerability, and report the vulnerability publicly after a specific period. Based on the vulnerabilities reported in the National Vulnerabilities Database (NVD), we estimate that market mechanisms accounted for 3.5 percent of all vulnerability disclosures in 2007, and 14 percent in 2008.

Despite their increased use, the impact of such markets is far from clear. There are two potential benefits of vulnerability markets: (1) increased incentives for benign security professionals to discover vulnerabilities, and (2) early but limited disclosure of vulnerability countermeasures to security service providers and other subscribers of the markets (in addition to the vendor) so that they can install defense mechanisms (such as detection signatures) before public disclosure of the vulnerability. If more systems are protected through limited disclosure, there will be less incentive for attackers to exploit the vulnerability even after public disclosure. Furthermore, by providing an alternative revenue source for attackers, some may choose to disclose through the markets rather than exploit the vulnerability. On the other hand, many have questioned the wisdom in creating incentives for discovering vulnerabilities and argue that (1) more discovered vulnerabilities will lead to more attacks, similar to weapons buy-back programs that can increase the number of guns on the street (Mullin 2001), (2) the presence of private intermediaries introduces additional complexity to the disclosure and patching cycle (Li and Rao 2007), (3) the private vulnerability markets have an incentive to leak vulnerability information to make their services more valuable (Kannan and Telang 2005), and (4) attackers may reverse engineer the attack signatures in the IDS provided by the markets to their clients and attack systems before a patch is available and deployed.

Therefore, the fundamental question remains open: Are markets for vulnerabilities effective in improving information security? We address one key aspect of this overall question through a large-scale empirical study that compares the impact of vulnerabilities disclosed through the market-based and nonmarket-based mechanisms. Specifically, we analyze over 2.4 billion information security alerts for 960 clients of an U.S.-based security service provider (SecureWorks) to examine three measures of impact: (1) attack diffusion (does market-based disclosure affect the diffusion of attacks corresponding to the vulnerability through the population of target systems), (2) attack risk (does market-based disclosure affect the likelihood of a vulnerability being exploited), and (3) attack volume (does market-based disclosure affect the volume of attacks that are based on the vulnerability). In addition, we examine vulnerability characteristics that increase the effectiveness of market-based disclosure. Thus, our research focuses on the impact of market-based vulnerability disclosure on the attack process and does not evaluate the discovery of vulnerabilities, the motivation of security researchers, or the types of vulnerabilities that are likely to be reported through the markets.

Recent research has examined optimal patching and vulnerability disclosure policies through analytical models (Arora et al. 2006; Arora et al. 2008; Cavusoglu et al. 2007). Some of the findings of this research are that it may be optimal for a software vendor to allow pirated copies to apply a patch (August and Tunca 2008), that it may be optimal for a monopolist to release software late with fewer vulnerabilities (Arora et al. 2006), that user rebates for good patching behavior can be optimal for proprietary software, but not for open source software (August and Tunca 2006), and that when vendors and users release and apply patches based on a time schedule, cost sharing can achieve social optimality (Cavusoglu et al. 2008). Research on optimal disclosure policies has found that vulnerability disclosure by agencies such as Computer Emergency Response Team (CERT) can force vendors to release patches more quickly (Arora et al. 2008), and that the risk of vulnerability exploitation, cost structure of the user population, and vendor's incentives to develop a patch determine the optimal disclosure policy (Cavusoglu et al. 2007). Related to our research, analytical models demonstrate that incentives to leak information make private vulnerability markets socially suboptimal (Kannan and Telang 2005).

There are three primary contributions of our research to the literature on optimal policies and methods to ensure the security of information systems. First, while several analytical models in the literature examine optimal vulnerability disclosure and patching policies, this research is one of a few that empirically evaluates a contemporary vulnerability disclosure phenomenon through the examination of IDS data, providing needed diversity in research methods (Mahmood et al. 2010;

Siponen 2005). Second, while economic models based on rational choice form the basis of the published research in this area (Arora et al. 2008; Cavusoglu et al. 2007; Kannan and Telang 2005), we develop our hypotheses through a review of the innovation diffusion literature (Rogers 2003; Van den Bulte 2000), providing additional diversity in the theoretical lenses used to study the phenomenon. More broadly, we contribute to the innovation diffusion literature by extending the concepts to a new, rich, and contemporary context that involves a negative innovation (the discovery of a software vulnerability); Rogers (2003, p. 106) emphasizes that "one of the most serious shortcomings of diffusion research is its pro-innovation bias." Finally, we empirically evaluate a research question that is of significant practical importance—whether vulnerability markets are effective. While such markets provide incentives for more benign discovery, it is not clear whether they actually improve the security of systems (Kannan and Telang 2005). We believe that our findings are of significant practical interest to policy makers and vendors.

The rest of the paper is organized as follows. We first provide a summary of the vulnerability disclosure process and the information security environment. In the subsequent section, we develop our hypotheses through a review of the innovation diffusion literature. We then describe the data used to empirically test our hypotheses. The next section contains the results of our analysis, and the final section concludes the paper.

The Information Security Environment

Vulnerability Patches and Countermeasures

Since software vulnerabilities cannot be completely eliminated in practice, the focus of information security has been on the effective *discovery* of vulnerabilities (Kannan and Telang 2005), on the rapid *patching* of systems to eliminate these vulnerabilities (August and Tunca 2008), and on the *countermeasures* to protect systems against these vulnerabilities such as installing intrusion detection systems (Cavusoglu et al. 2005) or disabling ports and functionality in affected software (Ransbotham and Mitra 2009). Agencies such as the Computer Emergency Response Team (CERT) play an important role in all stages of the security process by acting as intermediaries between security professionals who discover vulnerabilities, vendors who patch systems, and users who deploy the patches to protect against potential attacks. CERT does not pay security professionals for vulnerability discovery, and disclosure to CERT is voluntary. Many other similar agencies disclose vulnerabilities but do not pay for discovery, such as Security Focus, XForce, Secunia, BugTraq and Internet Security Systems X-Force. The National Vulnerabilities Database (NVD) consolidates the vulnerabilities reported in these separate data sources into a single database for easy access and tracking. Security service providers, such as Internet Security Systems (IBM) and SecureWorks, install intrusion detection devices in their client networks based on the vulnerability information they receive, monitor suspicious activity, and help clients install countermeasures and patches. Security vendors such as McAfee and Symantec provide security software and devices to companies, governments, and individuals.

There are two methods for removing or reducing the effect of software vulnerabilities. First, patches that are developed and distributed by the software vendor attempt to remove the vulnerability through an update to the software itself (Arora et al. 2006; Cavusoglu et al. 2008). Patches correct the vulnerability but can sometimes introduce other vulnerabilities in the software. Second, even if a patch is not available or installed, specific countermeasures can provide partial protection against attacks. Ransbotham and Mitra (2009) describe three types of countermeasures in systems that limit the impact of a vulnerability: (1) access control methods that limit access to the affected software to specific groups, (2) feature control methods that disable functionality and features in the affected software and devices, and (3) traffic control methods that filter suspicious traffic based on signature-based attack detection. Countermeasures are easier to develop than patches, but they provide temporary and imperfect protection until the vulnerability is removed through patching or a software upgrade. For some vulnerabilities, countermeasures may also be ineffective or impractical to implement.

Pathways to Vulnerability Disclosure

Figure 1 summarizes the primary pathways to public disclosure of vulnerabilities based on Frei et al. (2009). When security professionals discover a vulnerability, there are three pathways through which the vulnerability can be disclosed to the public. First, security professionals may choose to publicly disclose the vulnerability immediately through security mailing lists such as BugTraq, a pathway we refer to as *immediate disclosure*. When disclosed through mailing lists such as BugTraq, the vulnerability information is immediately disseminated to a wide audience of security professionals who can install countermeasures, to vendors who can

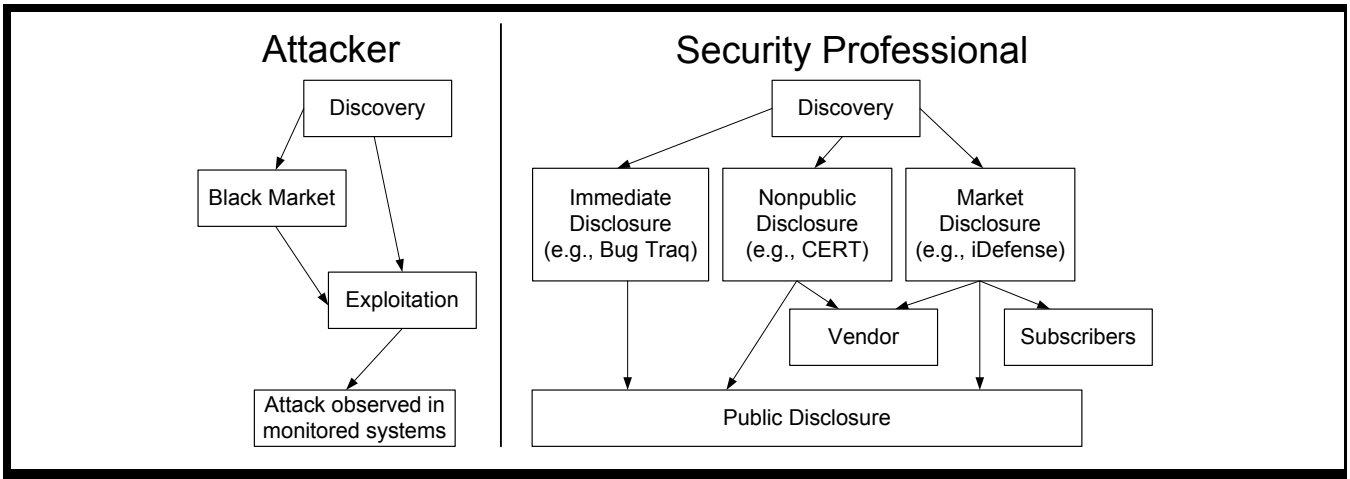

Figure 1. Pathways to Vulnerability Disclosure

develop patches, and to potential attackers who can exploit the information to their advantage. Immediate disclosure is often motivated by the need to force unresponsive vendors to address a vulnerability and to create incentives for developing secure software (Arora et al. 2006; Arora et al. 2008). The vendor, security professionals, and potential attackers receive notification of the vulnerability at the same time.

Second, the discoverer may choose to report the vulnerability to CERT, a pathway we refer to as *nonpublic disclosure*. CERT immediately notifies the vendor and discloses the vulnerability to the public when a patch is available from the vendor, or after a specific period (typically 45 days after notifying the vendor). In this disclosure path, the vendor is notified immediately after CERT receives the vulnerability details, while security service providers, security vendors, and potential attackers receive notification at the time of public disclosure. Even though CERT does not pay for reported vulnerabilities, nonpublic disclosure through CERT is the preferred disclosure path for many security professionals because it gives vendors a reasonable period to develop patches, while retaining incentives for the vendor to develop solutions in a timely manner (Arora et al. 2008).

Third, security professionals can sell the vulnerability to the vulnerability markets described earlier (iDefense and ZDI), a pathway we refer to as *market disclosure*. The markets pay an undisclosed amount that is based on the severity of the vulnerability, the popularity of the underlying product, and the exclusivity of the information. Recent surveys indicate that most payments for vulnerabilities are below $5,000, but can be as high as $30,000 for select vulnerabilities (Unsecurity Research 2010). The market providers inform the vendor after verifying the vulnerability so that the vendor can develop patches. A crucial difference with CERT is that the markets include the encrypted vulnerability signature in the intrusion detection systems (IDS) that they provide to their subscribers so that their subscribers have some protection even before the vendor develops patches to correct the vulnerability. The markets also issue advisories to their subscribers that do not disclose vulnerability details but include general countermeasures such as closing certain ports and limiting access to or disabling functionality in affected software (Tipping Point 2010). The subscribers for the markets (iDefense and Tipping Point) include managed security service providers, security product vendors and several Fortune 500 companies. Since many large and medium sized companies outsource security management to managed security service providers (Ransbotham and Mitra 2009), the encrypted vulnerability signatures and other countermeasures are widely disseminated. The market providers publicly disclose the vulnerability after a delay; for example, iDefense delays 180 days after vendor notification (iDefense Labs 2010).

On the other hand, when criminals who intend to exploit the vulnerability discover the vulnerability first, they may sell the vulnerability on the black market or they may exploit the vulnerability themselves. Anecdotal evidence suggests that prices obtained on the black market for critical vulnerabilities can be large, perhaps even as high as $120,000 (Frei et al. 2009). Security professionals discover exploited vulnerabilities only after monitoring systems detect attacks based on the vulnerabilities. Following discovery by security profes-

sionals, public disclosure of the vulnerability can take place through the three pathways described above.

Theory and Hypothesis Development

Diffusion of Attacks and Countermeasures

Once a vulnerability is discovered by attackers, the diffusion of attacks corresponding to the vulnerability through the population of target firms can be analyzed using concepts and patterns from the diffusion of innovations literature (Bass 1969; Rogers 2003; Van den Bulte 2000). The exploitation of a vulnerability by an attacker to attack a specific firm can be viewed as the adoption of an innovation (albeit a negative one). Similarly, the installation of protective measures (such as patches and other countermeasures) by target firms can also be viewed as a diffusion process. The innovation diffusion literature has developed models of the diffusion process (Bass 1969; Mahajan 1985; Rogers 2003) that have been extended over time to include a variety of additional factors such as the presence of key individuals and hubs (Goldenberg et al. 2009), learning by consumers (Lieberman 1987), and the social and cognitive gaps between the groups involved (Ferlie et al. 2005). From the perspective of the target firm, there is an inherent race between the diffusion of attacks and the diffusion of patches and countermeasures. A target firm is at risk if it does not install patches and countermeasures before its systems are attacked.

Figure 2 shows the diffusion process of attacks, patches and countermeasures. Both the exploitation of a vulnerability by attackers and the protection of systems by security professionals can be conceptualized as multistage processes (see Figure 2). The discovery of a vulnerability initiates an effort by attackers to develop methods to exploit the vulnerability (Kannan and Telang 2005). After an exploit method is developed, attacks are not instantaneous, but diffuse through the population of target systems over time as attackers discover more systems to attack (Ransbotham and Mitra 2009).

Likewise, the discovery of a vulnerability initiates an effort by the vendor to develop patches to remove the vulnerability (Cavusoglu et al. 2008). Once a patch is developed by the vendor, the vulnerability is not instantaneously eliminated, but patches also diffuse through the population of target systems over time (August and Tunca 2008). Similarly, the discovery of a vulnerability leads to the development of countermeasures that reduce the impact of the vulnerability by security professionals. Once a countermeasure is developed, it also diffuses through the population of target systems and is more likely to be installed in firms that subscribe to services from security service providers or the vulnerability markets.

Modeling the Diffusion Process for Attacks

Our primary interest in this paper is on the impact of market-based disclosure of vulnerabilities on the diffusion of attacks based on the vulnerability. We model the diffusion of attacks through the population of target systems through the familiar s-curve that has been extensively used to model the diffusion of innovations in the literature (Johansson 1979; Rogers 2003; Van den Bulte and Stremersch 2004) for several reasons. First, Ransbotham and Mitra (2009) emphasize the role of the attacker subculture (Sutherland 1947) in the diffusion of attacks through the population of target systems. After a vulnerability is discovered by the subculture, it is initially exploited by a few attackers who have the expertise required to take advantage of the vulnerability. Subsequently, the exploit is packaged into tools and disseminated, and a larger number of less expert attackers can also exploit the vulnerability. Over time, interest in the vulnerability decreases as more systems are protected and expert attackers move to other opportunities (Ransbotham and Mitra 2009). The above process suggests that the number of new systems attacked over time follows the familiar new product adoption pattern of Rogers (2003). That is, while there are few attacks initially, it increases to an intermediate peak value as more attackers have the interest and means to exploit the vulnerability, and it then decreases over time as interest wanes in the attacker community. Thus, the cumulative number of systems affected over time follows the well known s-curve that has been widely used in the diffusion of innovation literature. Second, the s-curve can capture two intuitive features of the attack diffusion process that are of interest to us—that attacks peak after a delay and that the number of affected systems asymptotically reaches a maximum penetration level. Third, we demonstrate in the results section that the s-curve provides a good fit for the attack diffusion process based on our empirical data.

To model the diffusion process of attacks, we introduce the following notation. Let $N(t)$ be the cumulative number of target systems affected at time t where t is measured from the time a method to exploit the vulnerability is discovered. Let P be the height of the s-curve, or the maximum number of target systems in the population affected by the vulnerability (referred to as penetration of the diffusion process). Let T_h be the time when half of the target systems are affected by the vulnerability. R is the slope of the s-curve which is dependent on factors such as the type of vulnerability, complexity of developing exploits, and the impact of the vulnerability on

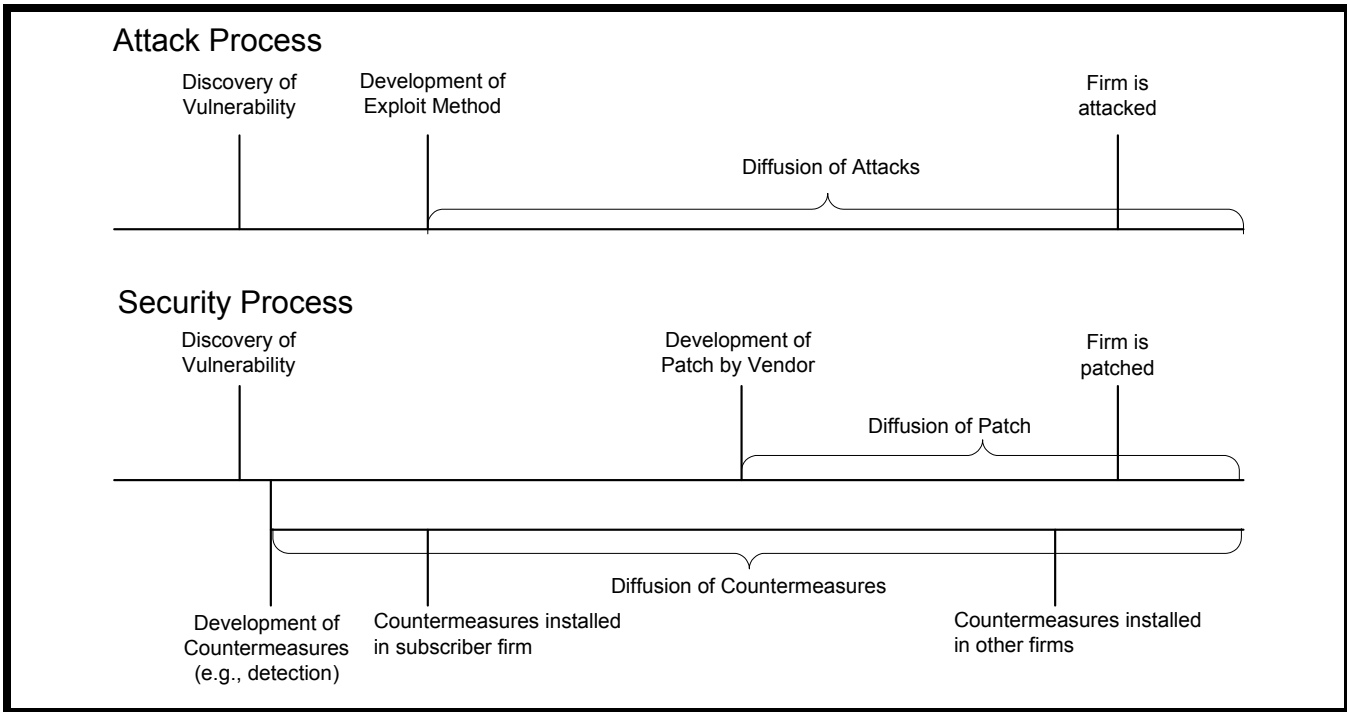

Figure 2. Diffusion of Attacks and Countermeasures

systems. $N(t)$ is modeled using the following familiar form of the s-curve:

$$N(t) = \frac{1}{1 + e^{-R(t-T_h)}} \quad (1)$$

Market Disclosure and the Diffusion Process

Market-based disclosure of vulnerabilities has two fundamental differences with the two other paths of disclosure (immediate disclosure and nonpublic disclosure) depicted in Figure 1. We describe these two differences below and they form the bases of the hypotheses that follow.

Limited Disclosure of Vulnerability Countermeasures

When a vulnerability is reported to market-based mechanisms like Tipping Point or iDefense, the encrypted signature is made available to subscribers of the market for immediate inclusion in their IDS. The markets also provide advisories for their subscribers that include countermeasures to minimize the effect of the vulnerability, such as temporarily closing certain ports, disabling functionality in software, or further limiting access to the affected software. Prior to public disclosure, the markets also provide details to security vendors who subscribe to their services so that they can better protect their clients (Tipping Point 2010). Since many large corporations outsource security management to security service providers and consumers are often protected through security software from security vendors, the signatures and countermeasures are widely disseminated. This limited disclosure of the vulnerability in the form of encrypted signatures and advisories to market subscribers ensures that a significant number of target systems would have the countermeasures deployed at the time of public disclosure. On the other hand, when vulnerabilities are disclosed through mailing lists such as BugTraq (immediate disclosure), security service providers, security professionals, vendors, and potential attackers obtain notification of the vulnerability at the same time, and countermeasures are not installed at the time of public disclosure. Even when vulnerabilities are disclosed through CERT, vendors are notified immediately but security service providers and security vendors receive notification of the vulnerability only at the time of public disclosure.

Ransbotham and Mitra (2009) describe information security compromise as a two-stage process where the attacker initially probes the population of target systems to find those that have the specific vulnerability and do not have counter-

measures installed, and then selectively attacks. Similar descriptions of the information security compromise process also appear in Kshetri (2005). For vulnerabilities disclosed through the market-based mechanisms, more systems are likely to have the countermeasures installed at the time the attacker attempts to exploit the vulnerability. Consequently, the target population of unprotected systems detected during the first stage of the attack process will be smaller, leading to fewer potential targets in the second stage, and a smaller height (P) of the s-curve for the diffusion process.

Hypothesis H1: The diffusion of attacks through the population of target systems will have reduced penetration for vulnerabilities reported through the market mechanisms.

Delayed Public Disclosure

When compared to immediate disclosure through mailing lists such as BugTraq, market disclosure of vulnerabilities introduces a time lag between vulnerability discovery and its public disclosure. We also argue that the time lag between initial vulnerability discovery and its public disclosure is greater for vulnerabilities disclosed through the market mechanisms when compared to vulnerabilities disclosed through CERT (nonpublic disclosure) for two reasons. First, private vulnerability markets work in their own best interests to increase the value of their services to their subscribers (Kannan and Telang 2005; Li and Rao 2007). Private information about a vulnerability is only valuable to the markets until everyone has access to the same information through public disclosure. As this period is extended, the services that the market provides to its subscribers (countermeasures against the vulnerability) are more valuable, and the incentives for the markets are aligned toward delayed public disclosure. The markets have multiple methods to delay public disclosure such as waiting longer to inform the vendor and giving the vendor a longer period to develop the patches. Second, unlike the private vulnerability markets (iDefense and Tipping Point), CERT has a public policy role and has to create incentives for the vendor to develop the patch in a timely manner (Arora et al. 2008). Consequently, the CERT vulnerability disclosure policy on their website explicitly states that vulnerabilities will be reported to the public after 45 days from initial report (CERT 2010). In contrast, the vulnerability disclosure policy for iDefense states that iDefense considers a reasonable timeframe to be no more than 180 days from initial report (iDefense Labs 2010).

The delayed public disclosure for vulnerabilities disclosed through the market introduces a time lag in the diffusion process of the vulnerability because attackers have reduced opportunity to become aware of the existence of the vulnerability until it is publicly disclosed. Thus, market disclosure delays the development of methods to exploit the vulnerability and the subsequent start of the attack diffusion process, and the s-curve in Equation (1) shifts to the right.

Hypothesis H2: The diffusion of attacks through the population of target systems will be delayed for vulnerabilities reported through the market mechanisms.

Markets also have an incentive to leak vulnerability information to increase the value of their services since their own clients are partially protected through installed countermeasures (Kannan and Telang 2005). Furthermore, in spite of precautions taken by the market providers, the encrypted signatures and the advisories they provide their clients can also be used by attackers to reverse engineer the vulnerabilities. Thus, alternatively, the diffusion of attacks could start earlier for market disclosed vulnerabilities as attackers capitalize on this information. However, we believe that the fear of negative publicity and potential legal consequence reduces incentives to leak vulnerability information and increases incentives to exercise caution when disclosing information.

Risk of First Attack

Early disclosure of countermeasures to market subscribers also reduces the likelihood of attack success since more target systems are likely have countermeasures installed by the time they are attacked. Rational choice predicts that experts focus on the more rewarding opportunities (Becker 1968; Ehrlich 1996), and the lower probability of success deters expert attackers who are typically the first to exploit a vulnerability (Ransbotham and Mitra 2009). Even if attackers are motivated by a need for recognition among peers, lower probability of attack success will make such vulnerabilities less attractive to expert attackers. Delayed public disclosure keeps the vulnerability away from the public domain for a longer period making it less visible to attackers, reducing the attention it receives from them, and shortening the period of successful exploitation before systems have countermeasures and patches installed. From the perspective of a target firm, this implies that they face a lower risk of first attack from expert attackers for vulnerabilities that are reported through the market mechanisms.

Hypothesis H3: A target firm will face a lower probability of first attack for vulnerabilities that are reported through the market mechanisms.

Volume of Attacks

While initial attacks may be from expert attackers, volume of attacks is primarily driven by the availability of automated tools that make the vulnerability exploitation methods more accessible (Jones 2007; Kshetri 2005; Radianti and Gonzalez 2007). The lower probability of success for vulnerabilities disclosed through the market mechanisms will affect the volume of attacks in two ways. First, it will generate less interest in the overall attacker community, resulting in lower volume of attacks. When attackers expect that a vulnerability has less likelihood of success, their best response is to divert their limited resources to vulnerabilities which they expect will have a greater likelihood of success. Second, the diffusion of innovation literature emphasizes the role of key individuals or social hubs in the diffusion process for new products (Goldenberg et al. 2009; Van den Bulte and Joshi 2007). Social hubs have a large number of network connections and greatly influence the subsequent adoption of the innovation in the general population (Goldenberg et al. 2009). In our context, attackers who package vulnerability exploitation methods into automated tools and raise awareness among the network of potential attackers through toolkits, postings, and discussions on the Internet play a vital role in the attack diffusion process. The lower probability of success for vulnerabilities disclosed through market mechanisms will generate less interest among such attackers to develop automated tools to exploit the vulnerability, also leading to lower attack volume. The above arguments are valid if attackers are driven by economic motives or by the need for recognition among peers in the attacker community.

> *Hypothesis H4: A target firm will have fewer attacks for vulnerabilities disclosed through the market mechanisms.*

Vulnerability Characteristics and the Effects of Market Disclosure

While the previous hypotheses highlight the positive impact of market-based disclosure, will the effectiveness of market disclosure vary based on the characteristics of a vulnerability? A primary benefit of market disclosure arises from the early disclosure of vulnerability countermeasures to market subscribers so that a significant segment of the population is partially protected at the time of public disclosure. We investigate two fundamental vulnerability characteristics that affect the efficacy of such countermeasures in protecting against the vulnerability: (1) the availability of a signature at the time of disclosure, and (2) the complexity of the process required to exploit a vulnerability.

Availability of Vulnerability Signature

Limited disclosure of vulnerability countermeasures to security service providers and other market subscribers is more effective when the signature of the vulnerability is available at the time of disclosure, so that security vendors and security service providers can immediately install the signatures in the intrusion detection systems they provide to their clients (Cavusoglu et al. 2005; Kemmerer and Vigna 2002). More systems with installed signatures for attack detection will lead to lower likelihood of attack success and consequently lower attack volume. Without the availability of a signature, vulnerability countermeasures are limited in scope and less effective because intrusion detection systems cannot easily identify attacks in progress (Cavusoglu et al. 2005; Kemmerer and Vigna 2002). Consequently, the impact of market disclosure on attack volume is greater for vulnerabilities that have a signature available at the time of disclosure.

> *Hypothesis H5: The effect of market disclosure on (a) the probability of first attack and (b) the volume of attacks will be greater for vulnerabilities that have a signature available when reported to the markets.*

While Hypothesis H5 predicts a moderating role of signature availability, its direct effect is harder to theorize. When disclosed through the nonmarket mechanisms, signatures are disclosed to security service providers, security vendors, and potential attackers at the same time. The availability of a signature will allow attackers to reverse engineer the vulnerability and develop attack methods leading to a greater volume of attacks. On the other hand, signature availability will enable security service providers to detect and deter attacks, leading a lower probability of success, and a lower attack volume.

Complexity of the Vulnerability Exploitation Process

Vulnerabilities also differ in the complexity of the process required to exploit the vulnerability. High complexity vulnerabilities require specialized access conditions such as elevated privileges, user actions (such as downloading files and e-mails), or social engineering techniques that can be detected by knowledgeable experts (Mell and Romanosky 2007). On the other hand, low complexity vulnerabilities do not require specialized access conditions or affect software products that must provide wide access to many users, the affected software configuration is widespread in use, and the vulnerability can be exploited remotely through relatively simple processes (Mell and Romanosky 2007).

Thus, countermeasures based on access control (Ransbotham and Mitra 2009) are less effective for low complexity vulnerabilities because such vulnerabilities do not require specialized access conditions, local access, or elevated access privileges, and the affected software must provide wide access to many users by design. Also, countermeasures based on feature control of software and devices (Ransbotham and Mitra 2009) are less effective for low complexity vulnerabilities because the affected configuration is widespread in use and limiting functionality or disabling ports may thus be impractical. On the other hand, attacks based on high complexity vulnerabilities can be made ineffective by countermeasures such as limiting access, disabling ports and features, or educating users (Mell and Romanosky 2007). Thus, market-based disclosure and the consequent early disclosure of countermeasures to market subscribers have a greater effect for high complexity vulnerabilities.

> *Hypothesis H6: The effect of market disclosure on (a) the probability of first attack and (b) the volume of attacks will be greater for vulnerabilities that require complex exploitation methods.*

Table 1 summarizes our hypotheses and the underlying theoretical arguments. Hypotheses H1 and H2 evaluate the impact of market-based disclosure on the attack diffusion process. Hypothesis H3 and H4 evaluate the impact of market-based disclosure on the risk of first attack and attack volume, respectively. Hypotheses H5 and H6 examine how the impact of market-based disclosure is affected by two fundamental vulnerability characteristics.

Data

To evaluate our hypotheses, we combine three main data sources: (1) a proprietary database of intrusion detection system alerts from a security service provider (SecureWorks), (2) publicly available information from two vulnerability markets (*iDefense* and *Tipping Point*) about the vulnerabilities that are reported through them, and (3) information available from the National Vulnerability Database (NVD) that combines several other public vulnerability data sources such as CERT, BugTraq, XForce, and Secunia. We believe that ours is the first study that combines the vulnerability market and NVD data with actual intrusion detection data from a large number of firms to empirically evaluate a contemporary information security issue.

Intrusion Alert Data

Our primary data source is a proprietary database of alerts generated from intrusion detection systems (IDS) installed in client firms of a security service provider (SecureWorks). Each time the IDS detects a signature in an incoming data stream, it generates an alert for further analysis. Cavusoglu et al. (2005) and Ransbotham and Mitra (2009) describe the role and function of intrusion detection systems in more detail. The dataset provides a unique research opportunity because it contains real alert data (as opposed to data from a research setting) from a large number of clients with varied infrastructure across many industries. The alert database contained over four hundred million alerts generated during 2006 and 2007. Our analysis is based on a panel dataset of the number of alerts generated every day during the two-year period of our analysis, summarized by target firm and specific vulnerability.

Market Disclosed Vulnerabilities

Our second data source is publicly available information from the existing vulnerability markets. During 2006 and 2007, there were two markets providing incentives for security researchers to discover and disclose vulnerabilities through their service (iDefense and Tipping Point). During that same period, there were many other options available for vulnerability disclosure that did not reward researchers directly, such as CERT, Security Focus, XForce, Secunia, BugTraq and Internet Security Systems X-Force. One key focal variable for us is an indicator variable, *Market*, that is set to 1 if a vulnerability was disclosed through one of the two market-based mechanisms, and 0 otherwise. Vulnerabilities disclosed through the market mechanisms are also subsequently reported to the NVD after a delay, and that database contains a wealth of additional detailed information about market and nonmarket-based vulnerabilities that we describe below.

National Vulnerabilities Database

Our third main data source is the National Vulnerabilities Database (NVD 2008). We match the signatures for each unique vulnerability in our intrusion alert database with detailed information available through the NVD. The matching is done through a CERT assigned unique identification that links all three of our databases together. Each vulnerability in the NVD is assessed by experts using a Common Vulnerability Scoring System (CVSS) (Mell and Romanosky 2007; Mell et al. 2006). The CVSS is an open, mature, and well-established (e.g., Frei et al. 2006; Jones 2007; Kotenko and Stepashkin 2006) definition of the fundamental characteristics of a vulnerability. While the scoring system has some shortcomings, it is objectively scrutinized by many interested parties and uniformly applied to all vulnerabilities. Details of

Table 1. Summary of the Hypotheses

Hypothesis Summary	Mechanism	Related References	Empirical Support
DIFFUSION OF ATTACKS			
H1: Will have reduced penetration for vulnerabilities reported through the markets	Early disclosure of countermeasures to security service providers and security vendors will reduce the population of unprotected target systems	(Kshetri 2005; Ransbotham and Mitra 2009; Rogers 2003)	H1 supported
H2: Will be delayed for vulnerabilities reported through markets	Incentives for the markets are aligned toward delayed public disclosure since their subscribers are already protected through installed countermeasures	(Arora et al. 2008; Kannan and Telang 2005; Li and Rao 2007; Rogers 2003)	H2 supported
RISK OF FIRST ATTACK			
H3: Is lower for vulnerabilities reported through the markets	Market disclosure reduces the likelihood of attack success and rational choice dictates that expert attackers move to more rewarding opportunities.	(Becker 1968; Ehrlich 1996; Ransbotham and Mitra 2009)	H3 supported
VOLUME OF ATTACKS			
H4: Is lower for vulnerabilities disclosed through markets	Reduced likelihood of success dampens interest in the attacker community and the development of automated tools, thereby reducing attack volume.	(Goldenberg et al. 2009; Van den Bulte and Joshi 2007)	H4 supported
H5: The effect of market disclosure on (a) the probability of first attack and (b) the volume of attacks will be greater for vulnerabilities that have a signature when reported to the markets	Availability of a signature increases the value of early disclosure of countermeasures to security service providers who install signatures in their IDS.	(Cavusoglu et al. 2005; Kemmerer and Vigna 2002)	H5(a) not supported H5(b) supported
H6: The effect of market disclosure on (a) the probability of first attack and (b) the volume of attacks will be greater for vulnerabilities that require complex exploitation methods	Market disclosure increases ability to deploy countermeasures; countermeasures are more effective for higher complexity vulnerabilities.	(Mell and Romanosky 2007; Ransbotham and Mitra 2009)	H6(a) not supported H6(b) partially supported

the CVSS scoring system used in our analysis are in Mell and Romanosky (2007). It is important for our analysis that we insure that the effects we see are due to market disclosure and not due to characteristics of the vulnerability itself. The uniform scoring system, along with other data in the NVD, provides us with several variables that we can utilize as controls.

In addition to the *Market* indicator variable, we use the following focal variables in our empirical analysis to evaluate the hypotheses derived from the information available through the NVD. Once the attacker has access, vulnerabilities require varying degrees of complexity to exploit and are categorized by experts as *low*, *medium*, or *high* complexity. We code *Low* complexity as the base level and include two indicator variables, *Med* and *High*. It is important to note that the complexity score is based on the difficulty of exploitation and not on the difficulty of detection and deterrence. We also include an indicator variable (*Sig*) that is set to 1 if a signature was available at the time that the vulnerability was disclosed to the public, 0 otherwise. Because disclosure through BugTraq is immediate, we include an additional variable (*BugTraq*) to capture the effects of immediate disclosure.

In addition to the focal variables described above, we use the following control variables derived from the information available through the NVD. First, the *impact* of a vulner-

ability is categorized by experts as affecting the disclosure of confidential information (*I_conf*), the integrity of data (*I_integrity*), or the availability of system resources (*I_avail*). A vulnerability can be coded by experts to have multiple potential impacts, and we use an indicator variable for each impact category that is set to 1 if the potential for the specific impact (confidentiality, integrity, and availability) is present, 0 otherwise. Second, the NVD classifies vulnerabilities into seven different *types* based on the specific software flaw that the vulnerability represents. These are (1) incorrect access privileges (*T_access*), (2) failure to handle incorrect input (*T_input*), (3) shortcomings in the design of software (*T_design*), (4) insufficient response to unexpected conditions (*T_exception*), (5) weak configuration of settings (*T_config*), (6) errors due to sequencing of events (*T_race*), and (7) uncategorized vulnerability types (*T_other*). Separate indicator variables are included for the first four types; however, there were insufficient configuration and race type vulnerabilities and these were grouped into the base type. We also include an indicator variable (*Patch*) that is set to 1 if a patch was available on the focal day of analysis, 0 otherwise. We also include the *Age* of the vulnerability (log transformed) at the time of our analysis, measured as the number of days since the vulnerability was reported.

Methods and Results

Diffusion of Attacks

For consistency of analysis, we excluded all client firms of the managed security service provider (MSSP) that had more than one intrusion detection device installed or that did not have data available for the entire study period. Table 2 presents descriptive statistics for 1,201 vulnerabilities defined in the intrusion detection system (IDS) of the MSSP for which we could match their alert signatures to the NVD database. In the alert database, attackers exploited only 333 (27 percent) of the 1,201 vulnerabilities in the MSSP database during the period of the study; the others are either older vulnerabilities included in the IDS for historical reasons or vulnerabilities that were never attacked. Of the exploited vulnerabilities, 140 were actively exploited and affected many firms in our sample. The 333 vulnerabilities were aligned with day 0 representing the date the vulnerability was reported to the market mechanisms or other reporting agencies. To evaluate hypotheses H1 and H2, for each of the 333 vulnerabilities and for each day in the research period subsequent to day 0, we calculated the cumulative number of firms that had experienced exploitation attempts corresponding to the vulnerability until that day, to build a panel with 132,768 observations. Since day 0 for all vulnerabilities is not on the same calendar date, vulnerabilities did not have the same number of observations. For robustness checks, we also performed the analysis with the 140 actively exploited vulnerabilities.

To assess our assumption that the diffusion of exploitation attempts follows an s-curve, we first visually examined the plots of the cumulative number of clients affected over time for a few randomly chosen vulnerabilities and found that the curves match the s-shape. More formally, we performed the following analysis to assess the fit of the observed diffusion pattern to an s-curve. For each of the 333 vulnerabilities in our sample, we used nonlinear regression to estimate the diffusion function in Equation (2) below with P (total penetration), R (rate of penetration), and D (delay before start of penetration) as estimated parameters.

$$N(t) = \frac{P}{1 + e^{-(Rt-D)}} \tag{2}$$

The purpose of this analysis was to examine whether the cumulative number of systems attacked over time ($N(t)$) could be modeled using an s-curve with appropriately chosen P, R, and D values. We evaluate the general fit of the s-curve shape to the observed data by examining the R^2 of each regression. When estimated individually, the nonlinear estimation of P and D fails to converge for vulnerabilities that are minimally exploited and thus have insufficient variation in $N(t)$ over time. For these minimally exploited vulnerabilities, the estimation procedure is unable to distinguish between almost zero penetration (low P) and extremely long delay (high D). For the others, the data fits an s-curve very well; the average R^2 of the nonlinear regressions that was 81.3 percent and the median R^2 was 99.3 percent. For comparison, we also investigated an alternative function ($N(t) = P - e^{-Rt}$) that also asymptotically reaches a penetration level P but does not have the s-shape, and found that the corresponding mean and median R^2 values were 25.9 percent and 16.2 percent, respectively.

Next, to examine the impact of market disclosure on the diffusion of attacks (P, R, and D values), we estimated Equation (2) with covariates using the attack diffusion data for 333 vulnerabilities and nonlinear least squares estimation. (For robustness, we also repeat the analysis for only the 140 actively exploited vulnerabilities.) In Equation (2), we now allow P, R, and D to be linear functions of our focal and control variables.

$$D = \beta_0^D + \beta_1^D Sig + \beta_2^D Med + \beta_3^D High + \beta_4^D Market + control\,variabes \tag{3}$$

Table 2. Sample Descriptive Statistics

		Market		Nonmarket	
Variable	**Value**	**Count**	**%**	**Count**	**%**
Access	Requires Local	21	13.13	117	11.24
	Requires Adjacent	4	2.50	10	0.96
	Network	135	84.38	914	87.80
Complexity	Low	73	45.63	542	52.07
	Medium	70	43.75	387	37.18
	High	17	10.63	112	10.76
Authentication	Not required	151	94.38	991	95.20
	Required	9	5.63	50	4.80
Confidentiality Impact	No	22	13.75	256	24.59
	Yes	138	86.25	785	75.41
Integrity Impact	No	17	10.63	242	23.25
	Yes	143	89.38	798	76.66
Availability Impact	No	14	8.75	189	18.16
	Yes	146	91.25	852	81.84
Vulnerability	Access	9	5.62	54	5.19
	Input	35	21.88	355	34.10
	Design	12	7.50	175	16.81
	Exception	12	7.50	104	9.99
	Environmental	0	0.000	2	0.19
	Configuration	1	.63	17	1.63
	Race Condition	4	2.50	11	1.06
	Other	1	0.63	16	1.54
Contains Signature	No	129	80.63	680	65.32
	Yes	31	19.38	361	2.72
Patch Available	No	43	26.88	501	48.13
	Yes	117	73.13	540	52.89

$$R = \beta_0^R + \beta_1^R Sig + \beta_2^R Med + \beta_3^R High + \beta_4^R Market + control\ variables \quad (4)$$

$$P = \beta_0^P + \beta_1^P Sig + \beta_2^P Med + \beta_3^P High + \beta_4^P Market + control\ variables \quad (5)$$

Note that in Equation (2), we replace D, R, and P with the linear function shown in Equations (3), (4), and (5), and we estimate the parameters in the resulting single equation through pooled nonlinear least squares estimation. The term ($R * T_h$) in Equation (1) is incorporated in the constant term in Equation (3). The data set for estimating Equation (2) contains for each vulnerability and for each date, the cumulative number of firms ($N(t)$) that is affected by the vulnerability until that date. The data also contains the values for the focal and control variables shown in Equations (3), (4), and (5). The focal and control variables change across vulnerabilities, and some variables (such as *Age*, *Sig*,and *Patch*) change across time as well (since a signature or patch may be made available at a later date). Hypotheses H1 and H2 predict that the two coefficients of the *Market* variable (β_4^D and β_4^P) are positive and negative, respectively.

Table 3 details the results of our analysis of vulnerability diffusion using pooled nonlinear least-squares estimation of parameters. Model 0 describes the effects of the control variables only on the overall penetration (P), rate of diffusion (R), and delay (D), and explains 18.8 percent of the variance in the diffusion patterns. Model 1 introduces our focal variables. We find support for Hypothesis H1 that market disclosure decreases the overall penetration of attacks (β = -73.362, $p < 0.001$ for the *Market* variable). We also find support for Hypothesis H2 that market disclosure delays the diffusion of exploitation attempts (β = 207.055, $p < 0.001$).

Table 3. Diffusion of Vulnerability Exploit Attempts

	Model 0			Model 1		
Variable	**P**	**R**	**D**	**P**	**R**	**D**
Constant	143.745*** (2.473)	-0.424*** (0.066)	286.396*** (44.660)	62.394*** (1.735)	-0.834*** (0.095)	53.485*** (6.380)
Confidentiality Impact (*I_conf*)	-94.762*** (2.200)	-0.791*** (0.124)	-15.307*** (2.658)	-37.478*** (1.500)	0.108*** (0.014)	106.275*** (11.986)
Integrity Impact (*I_integ*)	-36.192*** (1.870)	-0.305*** (0.050)	16.291*** (2.785)	9.933*** (1.553)	0.280*** (0.032)	75.372*** (8.530)
Availability Impact (*I_avail*)	10.451*** (1.802)	-0.045** (0.014)	-240.862*** (37.470)	-12.520*** (1.438)	-0.618*** (0.070)	-116.628*** (13.191)
Input Type (*T_input)*	25.407*** (0.774)	0.399*** (0.062)	-5.768*** (1.977)	58.148*** (0.895)	0.386*** (0.044)	94.796*** (10.584)
Design Type (*T_design*)	14.190*** (1.237)	0.139*** (0.022)	-5.768*** (1.365)	-31.989*** (1.357)	-0.289*** (0.033)	8.494*** (1.271)
Exception Type (*T_exception*)	141.466*** (2.530)	0.456*** (0.072)	12.906*** (2.044)	33.645*** (2.439)	-1.184*** (0.133)	488.845*** (55.231)
BugTraq Disclosure (*BugTraq)*	-12.184*** (0.764)	-0.297*** (0.046)	14.407*** (2.446)	3.885*** (0.833)	-0.062*** (0.008)	-5.392*** (0.870)
Patch Available (*Patch)*	40.797*** (0.781)	1.013*** (0.158)	-31.446*** (5.017)	-23.233*** (0.868)	-0.466*** (0.053)	-110.736*** (12.475)
Medium Complexity (*Med*)				212.146*** (1.823 7)	0.439*** (0.050)	106.987*** (11.948)
High Complexity (*High*)				39.710*** (1.117)	0.078*** (0.010)	17.079*** (2.152)
Signature Available (*Sig*)				143.019*** (1.147)	1.133*** (0.127)	-112.902*** (12.849)
Market Disclosure (*Market*)				-73.362*** (1.496)	-0.876*** (0.098)	207.055*** (23.349)
R^2			18.55%			32.14%

132,768 daily observations of 333 vulnerabilities from 2006–2007

Robust standard errors in parentheses; significance: *p < .05; **p < .01; ***p < .001

Nonlinear regression on number of firms affected, $N(t) = \frac{P}{1 + e^{-(4t-D)}}$ where the cumulative penetration (*P*), the rate of diffusion (*R*) and delay (*D*) are linear functions of the variables shown in the table.

When only the 140 actively exploited vulnerabilities are included in the analysis, the results are similar and support the hypotheses.

The nonlinear curve makes interpretation of coefficient parameters challenging in Table 3. To ease interpretation of estimates, Figure 3 graphs the resultant curves and illustrates the differences due to market-based disclosure. To generate the figure, we take the parameter estimates from Table 3 and use them in Equation (2). For the covariates other than *Market*, we use their mean value within the focal subset (market or nonmarket, as appropriate). Through the resultant curves, we see both the delay in diffusion and overall reduction in penetration for market disclosed vulnerabilities. The graph supports Hypothesis H1 and Hypothesis H2. They illustrate that market-based disclosure delays the diffusion process by approximately 30 days, and reduces penetration from 12 percent of firms to 9.5 percent of firms (a 21 percent decrease). Interestingly, once diffusion starts, the diffusion is rapid with the curves almost vertical. The rapid diffusion underscores the importance of additional time to implement countermeasures.

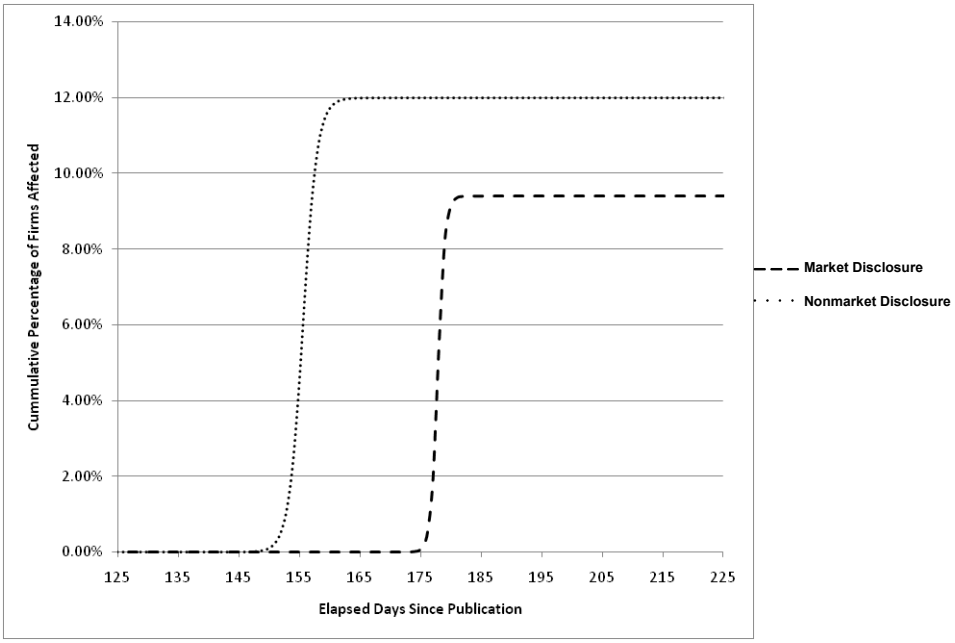

Figure 3. Diffusion of Vulnerability Exploit Attempts

The use of time-invariant explanatory variables in a pooled analysis can lead to biased estimates (Plumper and Troeger 2007). Thus, as a robustness check, we also performed an alternative two-step procedure.[2] We first estimated the *P*, *R*, and *D* values for each vulnerability individually through nonlinear least squares estimation. For market-based vulnerabilities, the average estimated value of penetration (*P*) is 36.5 (out of 960 firms) and the average estimated delay (*D*) is 186.0 days. For nonmarket vulnerabilities, the average estimated value of penetration (*P*) is 141.4 firms and the average estimated delay (*D*) is 79.9. These differences are practically and statistically significant based on a simple t-test, and consistent with Hypotheses H1 and H2. We also use the predicted values of penetration (*P*) and delay (*D*) for each vulnerability as the dependent variable in Equations (3) and (5) to estimate the two coefficients of the *Market* variable (β_4^D and β_4^P). The dependent variables in the second stage are the results of a prior nonlinear regression and thus not only have estimates of their values, but also associated standard errors. These standard errors reflect the accuracy of the prior nonlinear regression. To correct for known problems that arise when using estimated values as dependent variables in a regression equation (Saxonhouse 1976), we use the inverse of the standard errors as weights in a weighted least squares estimation to assess the effect of market-based disclosure on penetration (*P*) and delay (*D*). (Based on post regression analysis of variance inflation factors, we limit the weighting of observations to 30; increased weighting introduces colinearity based on the variance inflation factors.) The regression results indicated that market-based disclosure significantly reduces the vulnerability penetration (β = -154.71, $p < 0.01$) and increases the delay (β = 146.70, $p < 0.01$). Overall, the results of this analysis are consistent with the analysis reported in Table 3. However, we focus our attention on the results in Table 3 instead of this alternative approach. When estimated individually, the nonlinear estimation fails to converge for those vulnerabilities that penetrated few firms and thus had insufficient variation in $N(t)$ over time. Furthermore, the weighting of observations in the second stage to correct for the estimation errors in *P* and *D* values is imperfect and introduces noise in the analysis (Lewis and Linzer 2005).

Risk of First Attack

We examine the risk of first attack from a vulnerability through a Cox (1972) proportional hazard model. In the hazard model, the event being explained is the first exploitation attempt of a vulnerability for a specific client. The hazard model allows us to integrate information for vulnerabilities in our data set that were never exploited, to incorporate the

[2]We are grateful to an anonymous reviewer for suggesting this alternative approach.

evolving risk of exploitation, and to handle truncation of observations caused by the end of the study period. To perform this analysis, we construct a data set that contains, for each firm and vulnerability combination, the day of first attempt to exploit the vulnerability (960 firms and 1,201 vulnerabilities for a total of 1,152,406 observations). As before, all vulnerabilities were aligned so that day 0 represented the date the vulnerability was reported to the market mechanisms or other reporting agencies. The Cox proportional hazard model consists of two parts: the baseline hazard function describing how the risk of first attack changes over time when all covariates are at the mean level, and a parameter for each covariate that describes how the baseline hazard changes in response to explanatory covariates. In the Cox model, the baseline hazard is not affected by the covariates, and the parameters are assumed to have a multiplicative effect on the baseline hazard. To incorporate unobserved heterogeneity in firm-specific risk, we stratify the analysis by firm so that the baseline hazard function in the Cox proportional hazard model is allowed to vary by firm. Furthermore, to evaluate Hypothesis H3, we incorporate the focal and control variables described earlier as explanatory variables in the hazard model.

Table 4 details the results of the proportional hazard analysis. Model 0 describes the effects of only the control variables on the risk of exploitation (baseline hazard). We see little explanatory power in this model. However, Model 1 enters our focal variables and supports Hypothesis H3. The coefficient of the *Market* variable (β = -1.599, p < 0. 001) indicates that the risk of first attack decreases significantly for vulnerabilities disclosed through the markets. Based on the estimated parameters, market disclosure reduces the likelihood of first attack by 80 percent ($e^{-1.599}$ = 0.20), an economically significant decrease.

However, contrary to Hypotheses H5(a) and H6(a), the coefficients of the *Market * Sig, Market * Med*, and *Market * High* variables in Table 3 are not significant. That is, while market disclosure significantly lowers the risk of first attack, the impact does not vary based on the availability of vulnerability signature or the complexity of a vulnerability. A possible explanation is that while delayed public disclosure of a market-reported vulnerability reduces its visibility and the risk of first attack, countermeasures are less relevant against expert attackers.

Post regression tests find no indication of model misspecification. Despite the large sample size, the link test of squared linear predictor significance does not indicate specification problems; the linear term is significant (β = 35.90, p < 0.001) while the squared linear term is insignificant (β = 1.42, p > 0.156). Tests of proportional hazard assumptions based on Schoenfeld and scaled Schoenfeld residuals (Grambsch and Therneau 1994) find no significant time varying coefficients; however, because we use stratified models to account for firm-specific differences in the baseline hazard, tests of proportional hazard assumptions are not conclusive.

Volume of Attacks

To evaluate Hypotheses H4–H6, we use a two-stage Heckman (1979) model analyze the number of alerts generated by a vulnerability for a specific firm. We construct a data set that has for each firm (960 firms) and each vulnerability (1,201 vulnerabilities), the number of alerts generated on each day of our research period. As before, we align all vulnerabilities so that day 0 represents the day the vulnerability was first reported to the markets or other agencies. Since day 0 for all vulnerabilities is not on the same calendar date, some vulnerabilities will have longer periods for which data is available, and the total number of observations in our panel data set is 1,302,931. In addition, many vulnerabilities are never exploited in our alert data and ordinary least squares estimation will ignore the selection bias.

The two-stage Heckman model allows us to incorporate selection bias in the volume of attacks. In the first stage, we use a selection model to investigate vulnerability attributes that affect overall likelihood of exploitation. We control for all vulnerability covariates and include monthly fixed effects based on vulnerability disclosure date to control for possible changes in exploitation propensity over the two-year sample period. In the second stage, we examine the number of alerts per day (with a natural log transformation). In this analysis, we control for all vulnerability covariates and we include monthly fixed effects based on attack date to control for changes in attack behavior over time. We also include 960 firm fixed effect indicators to control for potential differences in a firm's inherent risk of attack. In the two models below, i indexes a firm, k indexes a vulnerability, and E_{ik} = 1if vulnerability k is ever exploited in our alert data, 0 otherwise. V_{ikt} is the volume of attacks on firm i on day t for vulnerability k. F_i^1 and F_i^2 are firm fixed effect dummies, while M_k^1 in stage 1 is the month fixed effects based on vulnerability publication date, and M_t^2 in stage 2 is the month fixed effect based on the attack date.

$$\text{1st Stage: } E_{ik} = \propto^{1} + \beta_1^1 Sig_k + \beta_2^1 Med_k + \beta_3^1 High_k + \beta_4^1 Market_k + \beta_5^1 Ln(Age_k) + M_k^1 + control\ variables \quad (6)$$

Table 4. Risk of Exploitation of Vulnerabilities

Variable	Model 0	Model 1	Model 2
Confidentiality Impact (*I_conf*)	0.028 (0.434)	-0.144 (0.442)	-0.141 (0.442)
Integrity Impact (*I_integ*)	0.266 (0.480)	0.315 (0.493)	0.313 (0.495)
Availability Impact (*I_avail*)	0.382 (0.409)	0.332 (0.422)	0.334 (0.421)
Access Type (*T_access*)	-0.382 (0.972)	-0.061 (0.974)	-0.058 (0.975)
Input Type (*T_input)*	0.292 (0.347)	0.292 (0.352)	0.286 (0.352)
Design Type (*T_design*)	-0.276 (0.566)	-0.353 (0.561)	-0.353 (0.562)
Exception Type (*T_exception*)	0.104 (0.573)	-0.113 (0.549)	-0.107 (0.550)
BugTraq Disclosure (*BugTraq)*	0.398 (0.316)	0.493 (0.315)	0.488 (0.316)
Patch Available (*Patch)*	0.040 (0.305)	-0.002 (0.291)	-0.004 (0.291)
Medium Complexity (*Med*)		-0.185 (0.436)	-0.185 (0.450)
High Complexity (*High*)		0.216 (0.357)	0.195 (0.369)
Signature Available (*Sig*)		1.085** (0.376)	1.064** (0.394)
Market Disclosure (*Market*)		-1.599*** (0.440)	-2.074*** (0.616)
*Market*Med*			0.625 (1.126)
*Market*High*			0.111 (1.093)
*Market*Sig*			0.467 (0.884)
Log likelihood	-114166.04	-111274.68	-111232.64
Wald χ^2	8.09	35.70***	50.23***

Cox proportional hazard model stratified by firm for 17,416 failures in 1,152,406 observations of 1,201 vulnerabilities and 960 firms; standard errors (in parentheses) are clustered by vulnerability.
Significance: *p < 0.05); **p < 0.01; ***p < 0.001.

$$\text{2ndStage: } Ln(V_{ikt}) = \propto^2 + \beta_1^2 Sig_k + \beta_2^2 Med_k + \beta_3^2 High_k + \beta_4^2 Market_k + \beta_5^2 Ln(Age_k) + \beta_6^2 Sig_k * Market_k + \beta_7^2 Med_k * Market_k + \beta_7^2 High_k * Market_k + F_i^2 + M_t^2 + control\ variables \quad (7)$$

Table 5 describes the results from the two-stage Heckman analysis. In stage 2 (Model 1), the coefficient of the *Market* variable is significant and negative (β = -0.116, p < 0.001), indicating a reduction in alert volume for vulnerabilities disclosed through the market by approximately 11 percent ($e^{-0.116}$ = 0.89). Thus, our results support Hypothesis H4. The results from stage 1 (Model 1) also provide additional support for Hypothesis H3, and indicate that market-disclosed vulnerabilities have a decreased likelihood of exploitation (β = -0.044, p < 0.001).

Table 5 (Model 2) shows the results when the interaction terms are introduced in the analysis. The coefficient for the *Market * Sig* interaction term is negative and significant (β = -0.097, p < 0.001), indicating that the negative impact of market disclosure on the volume of alerts is greater for vulnerabilities that have a signature available at the time of disclosure. Thus, our results support Hypothesis H5(b).

Table 5 (Model 2) shows that Hypothesis H6(b) is supported for medium complexity, but not for high complexity vulnerabilities. The coefficient for *Market * Med* is negative and significant (β = -0.065, p < 0.001) indicating that the negative impact of market disclosure on volume of alerts is greater for medium complexity vulnerabilities when compared to low complexity vulnerabilities. However, the coefficient for *Market * High* is positive and marginally significant. Thus, we find only limited support for Hypothesis H6(b).

Analysis of Selected Control Variables

Although not included in our hypotheses that focus on the impact of market-based disclosure, some control variables in our models also provide interesting insights. Of particular interest is the direct effect of the *Sig* variable in various models that indicate the availability of vulnerability signature at the time of disclosure. In Table 3 (Model 1), the availability of a signature reduces diffusion delay (β = -112.902, p < 0.001) and increases diffusion penetration (β = 143.019, p < 0.001). In Table 4 (Model 1), signature availability increases the risk of first attack (β = 1.085, p < 0.005), while in Table 5 (Model 1) signature availability at the time of disclosure increases alert volume (β = 0.102, p < 0.001). Clearly, our results indicate that attackers may learn to exploit a vulnerability by reverse engineering the associated signatures, thereby reducing diffusion delay and increasing diffusion penetration, attack risk, and attack volume.

Immediate vulnerability disclosure through BugTraq has only a marginal effect on attacks based on the vulnerability. In Table 3 (Model 1), the coefficient for the *BugTraq* variable is negative and significant (β = -5.392, p < 0.001 in the column labeled *D*) indicating that immediate disclosure through BugTraq reduces diffusion delay as expected. However, when compared to the constant term (β = 53.485, p < 0.001 in the column labeled *D*) the effect of BugTraq disclosure on diffusion delay is small. Likewise, BugTraq disclosure increases attack penetration (β = 3.885, p < 0.001 in the column labeled *P*) but the effect is small when compared to the constant term (β = 62.394, p < 0.001). Based on Tables 4 and 5, BugTraq disclosure does not affect risk of first attack, and only slightly decreases attack volume (β = -0.006, p < 0.005 for the BugTraq variable in Model 1 Stage 2 of Table 5).

Overall, the availability of a patch at the time of disclosure has a marginal effect on attacks based on the vulnerability. In Table 3 (Model 1), the coefficient of the *Patch* variable is negative and significant (β = -23.233, p < 0.001 in the column labeled *P*), indicating that the availability of a patch at the time of disclosure reduces the penetration of attacks based on the vulnerability as more systems have the vulnerability removed. However, Table 4 (Model 2) and Table 5 (Model 2) indicate that the availability of a patch at the time of public disclosure has an insignificant effect on attack risk and attack volume (the coefficient of the *Patch* variable is not significant in Table 4 Model 1 and small but marginally significant in Table 5 Model 1).

Summary and Conclusions

Our empirical results highlight that market disclosure of vulnerabilities (1) delays the onset and reduces the penetration of the attack diffusion process, (2) decreases the risk of first attack, and (3) decreases the volume of attacks corresponding to a vulnerability. We argue that the underlying reasons behind our empirical results are the delayed public disclosure of market reported vulnerabilities and the early but limited disclosure of vulnerability countermeasures to security service providers and security vendors. Furthermore, the impact of market disclosure on attack volume is greater for those vulnerabilities that have a signature available when reported to the markets. We find only limited evidence that the impact of market disclosure on attack volume is greater for complex vulnerabilities. The last column in Table 1 summarizes the empirical support for the hypotheses.

Table 5. Volume of Alerts per Client Firm per Vulnerability

Variable	Model 0		Model 1		Model 2	
Second Stage: Natural log of number of alerts						
Constant	0.546***	(0.081)	0.515***	(0.081)	0.513***	(0.081)
Confidentiality Impact (*I_conf*)	0.068***	(0.003)	0.059***	(0.003)	0.057***	(0.003)
Integrity Impact (*I_integ*)	-0.163***	(0.004)	-0.113***	(0.004)	-0.115***	(0.004)
Availability Impact (*I_avail*)	0.022***	(0.003)	-0.015***	(0.003)	-0.012***	(0.003)
Access Type (*T_access*)	-0.666***	(0.005)	-0.646***	(0.006)	-0.642***	(0.006)
Input Type (*T_input)*	0.086***	(0.002)	0.092***	(0.002)	0.094***	(0.002)
Design Type (*T_design*)	-0.043***	(0.003)	-0.078***	(0.003)	-0.079***	(0.003)
Exception Type (*T_exception*)	-0.128***	(0.004)	-0.143***	(0.004)	-0.146***	(0.004)
Age (ln)	-0.165***	(0.002)	-0.187***	(0.002)	-0.186***	(0.002)
BugTraq Disclosure (*BugTraq)*	-0.003	(0.002)	-0.006**	(0.002)	-0.005**	(0.002)
Patch Available (*Patch)*	0.012***	(0.002)	0.004**	(0.002)	0.002	(0.002)
Attack Month fixed effects	Included		Included		Included	
Firm fixed effects	Included		Included		Included	
Medium Complexity (*Med*)			-0.062***	(0.003)	-0.055***	(0.003)
High Complexity (*High*)			-0.032***	(0.003)	-0.030***	(0.003)
Signature Available (*Sig*)			0.102***	(0.003)	0.110***	(0.003)
Market Disclosure (*Market*)			-0.116***	(0.003)	-0.057***	(0.006)
*Market*Med*					-0.065***	(0.007)
*Market*High*					0.018*	(0.011)
*Market*Sig*					-0.097***	(0.009)
Inverse Mills	-0.151***	(0.004)	-0.089***	(0.004)	-0.091***	(0.004)
First Stage: Uncensored if vulnerability exploit attempt is ever seen in the sample						
Constant	0.143***	(0.008)	0.143***	(0.008)	0.143***	(0.008)
Confidentiality Impact (*I_conf*)	0.020***	(0.004)	0.020***	(0.004)	0.020***	(0.004)
Integrity Impact (*I_integ*)	0.511***	(0.004)	0.511***	(0.004)	0.511***	(0.004)
Availability Impact (*I_avail*)	-0.279***	(0.004)	-0.279***	(0.004)	-0.279***	(0.004)
Access Type (*T_access*)	-0.235***	(0.006)	-0.235***	(0.006)	-0.235***	(0.006)
Input Type (*T_input)*	0.134***	(0.003)	0.134***	(0.003)	0.134***	(0.003)
Design Type (*T_design*)	-0.202***	(0.004)	-0.202***	(0.004)	-0.202***	(0.004)
Exception Type (*T_exception*)	0.542***	(0.006)	0.542***	(0.006)	0.542***	(0.006)
BugTraq Disclosure (*BugTraq)*	-0.042***	(0.003)	-0.042***	(0.003)	-0.042***	(0.003)
Patch Available (*Patch)*	-0.448***	(0.003)	-0.448***	(0.003)	-0.448***	(0.003)
Publication Month fixed effects	Included		Included		Included	
Medium Complexity (*Med*)	0.052***	(0.003)	0.052***	(0.003)	0.052***	(0.003)
High Complexity (*High*)	0.269***	(0.004)	0.269***	(0.004)	0.269***	(0.004)
Signature Available (*Sig*)	0.759***	(0.004)	0.759***	(0.004)	0.759***	(0.004)
Market Disclosure (*Market*)	-0.044***	(0.004)	-0.044***	(0.004)	-0.044***	(0.004)
Wald χ^2	2.21e+06***		2.21e+06***		2.21e+06***	

Heckman two stage regression; n = 1,302,931; 709,090 uncensored; 343 vulnerabilities; robust standard errors in parenthesis. Two-tailed significance: *p < 0.05; **p < 0.01; ***p < 0.001

Are Markets Effective?

We find that the answer based on our analysis is nuanced. Consider that the life cycle of a vulnerability consists of four nonsequential stages: discovery of the vulnerability by attackers and security professionals, development of corrective measures by the vendor, deployment of corrective measures by the user community, and exploitation of the vulnerability by attackers. In this research, we focused on the exploitation of vulnerabilities by attackers. We observe from our empirical analysis that market disclosure delays the onset and reduces the penetration of attack diffusion, reduces the risk of first attack, and reduces the volume of attacks. Thus, if all vulnerabilities were indeed disclosed through the markets, there would be fewer attacks and the security environment would improve. Furthermore, one can argue that it is not only subscribers of the markets who benefit from the early disclosure of vulnerability countermeasures, but these benefits also accrue to others as overall attack volume and risk decreases.

However, as attacker resources are diverted from market-based vulnerabilities, they may be diverted to nonmarket-based vulnerabilities. While prior research has been concerned with information leakage from markets (Kannan and Telang 2005), there may be an equally important diversionary concern that we did not investigate here. In addition, preliminary analysis shows that there may be significant differences between market and nonmarket disclosed vulnerabilities in terms of complexity, authentication requirements, type of access required to exploit, and the impact of the vulnerability. If only certain types of vulnerabilities are reported through the markets (e.g., less complex vulnerabilities with local access requirements), then the impact of the markets will be limited. Furthermore, incentives for responsible disclosure will affect the discovery stage—a stage our research does not examine. In summary, while we find some beneficial effects of market-based disclosure on attack diffusion, attack risk, and attack volume, more research is needed that examines all stages in the life cycle of vulnerabilities.

Improving the Reporting of Vulnerabilities

Our findings suggest that some of the benefits of the market-based mechanisms over disclosure through industry organizations such as CERT accrue from the limited and early disclosure of vulnerability countermeasures to security service providers and security vendors. Industry organizations such as CERT could also establish the limited and early disclosure mechanisms, in addition to the current third party and private mechanisms that exist. At the time a vulnerability is discovered and reported to CERT, it can be disclosed to trusted security organizations so that the signature is deployed immediately in their IDS and other countermeasures are implemented for their clients. CERT can implement safeguards to ensure that this information is not disclosed to potential attackers by verifying the credentials of security organizations, encrypting the signatures provided, and only disclosing countermeasures and not the details of the actual vulnerabilities. Since many large corporations are protected by managed security service providers and many consumers are protected through software from security vendors, early disclosure of countermeasures will protect a large number of vulnerable systems.

Vulnerability markets implemented by CERT will have several potential benefits. Many security professionals do not report vulnerabilities to the private markets, thereby reducing the efficacy of the market mechanisms. Furthermore, private markets have incentives for information leakage or otherwise acting in their best interests (Kannan and Telang 2005), thereby reducing social welfare. CERT can recover the costs associated with paying vulnerability contributors through the membership fees it can charge to security organizations. Additionally, with competition between multiple markets (CERT, iDefense, and Tipping Point), vulnerability contributors will receive a higher price, thereby improving the discovery process.

Interestingly, in recent months, a few software vendors initiated programs to pay security researchers for vulnerabilities reported to them. For example, Google announced that the company will increase the maximum amount it pays for reported vulnerabilities in the Chrome operating system to approximately $3,000 (Fisher 2010). While the amount paid is still relatively small, it is a step toward motivating responsible discovery and disclosure by security professionals. Recently, other "white hat" vulnerability markets have also emerged. For example, the Exploit Hub is an online market where security researchers can sell validated exploits to known security professionals who can use the information to protect systems (Lemos 2010). Furthermore, while Microsoft does not pay for vulnerability disclosure and discourages public disclosure of vulnerability information, the company's recently announced *Coordinated Disclosure Policy* recognizes the need for early (but limited) disclosure to security professionals under certain circumstances (Thomlinson 2010). Thus, although still preliminary, there is growing interest in the concepts around market-based disclosure and limited disclosure to trusted security professionals, and empirical research can provide valuable insights.

Limitations and Future Research

There are several limitations of this study that future research can address. First, while we are able to examine the exploitation of vulnerabilities, the effects on the discovery process are unobserved. It will be a difficult (albeit important) task to assess the effects of the incentive mechanisms on market disclosure and future research could identify and evaluate optimal incentive schemes. Second, while the IDS and NVD data used in this research is voluminous and rich, it has several limitations. IDS data is inherently noisy and error-prone due to the large number of false positives and false negatives, and the categorization of vulnerabilities in the NVD database is often subjective. Better statistical methods to analyze this data set will enable researchers to empirically examine a wide range of research questions and to validate some of the analytical findings. There is a paucity of empirical research in this area that the analysis of IDS data can partially correct.

References

Anderson, R., and Moore, T. 2006. "The Economics of Information Security," *Science* (314:5799), pp. 610-613.

Arora, A., Caulkins, J. P., and Telang, R. 2006. "Sell First, Fix Later: Impact of Patching on Software Quality," *Management Science* (52:3), pp. 465-471.

Arora, A., Telang, R., and Hao, X. 2008. "Optimal Policy for Software Vulnerability Disclosure," *Management Science* (54:4), pp. 642-656.

August, T., and Tunca, T. I. 2006. "Network Software Security and User Incentives," *Management Science* (52:11), pp. 1703-1720.

August, T., and Tunca, T. I. 2008. "Let the Pirates Patch? An Economic Analysis of Software Security Patch Restrictions," *Information Systems Research* (19:1), pp. 48-70.

Baskerville, R. 1993. "Information Systems Security Design Methods: Implications for Information Systems Development," *ACM Computing Surveys* (25:4), pp. 375-414.

Bass, F. M. 1969. "A New Product Growth Model for Consumer Durables," *Management Science* (15:5), pp. 215-227.

Becker, G. 1968. "Crime and Punishment: An Economic Approach," *Journal of Political Economy* (76:2), pp. 169-217.

Cavusoglu, H., Cavusoglu, H., and Raghunathan, S. 2007. "Efficiency of Vulnerability Disclosure Mechanisms to Disseminate Vulnerability Knowledge," *IEEE Transactions on Software Engineering* (33:3), pp. 171-185.

Cavusoglu, H., Cavusoglu, H., and Zhang, J. 2008. "Security Patch Management: Share the Burden or Share the Damage?," *Management Science* (54:4), pp. 657-670.

Cavusoglu, H., Mishra, B., and Raghunathan, S. 2004. "The Impact of Internet Security Breach Announcements on Market Value of Breached Firms and Internet Security Developers," *International Journal of Electronic Commerce* (9:1), pp. 69-104.

Cavusoglu, H., Mishra, B., and Raghunathan, S. 2010. "The Value of Intrusion Detection Systems in Information Technology Security Architecture," *Information Systems Research* (16:1), pp. 28-46.

CERT. 2010. "CERT/CC Vulnerability Disclosure Policy," Software Engineering Institute, Carnegie Mellon University (http://www.cert.org/kb/vul_disclosure.html).

Cox, D. R. 1972. "Regression Models and Life Tables," *Journal of the Royal Statistical Society* (34:2), pp. 187-202.

Dhillon, G., and Backhouse, J. 2001. "Current Directions in IS Security Research: Towards Socio-Organizational Perspectives," *Information Systems Journal* (11:2), pp. 127-153.

Ehrlich, I. 1996. "Crime, Punishment and the Market for Offences," *Journal of Economic Perspectives* (10:1), pp. 43-67.

Ferlie, E., Fitzgerald, L., Wood, M., and Hawkins, C. 2005. "The Nonspread of Innovations: The Mediating Role of Professionals," *Academy of Management Journal* (48:1), pp. 117-134.

Fisher, D. 2010. "Google Fixes 10 Bugs in New Chrome Release," Post, The Kaspersky Lab Security News Service, September 15 (www.threatpost.com).

Frei, S., May, M., Fiedler, U., and Plattner, B. 2006. "Large-Scale Vulnerability Analysis," in *Proceedings of the 2006 SIGCOMM Workshop on Large-Scale Attack Defense*, New York: ACM Press, pp. 131-138.

Frei, S., Schatzmann, D., Plattner, B., and Trammell, B. 2009. "Modeling the Security Ecosystem–The Dynamics of (In)Security," in *Workshop on the Economics of Information Security*, University College, London, England (http://weis09.infosecon.net/files/103/paper103.pdf).

Goldenberg, J., Han, S., Lemann, D. R., and Hong, J. W. 2009. "The Role of Hubs in the Adoption Process," *Journal of Marketing* (73), pp. 1-13.

Gordon, L. A., and Loeb, M. P. 2002. "The Economics of Information Security Investment," *ACM Transactions on Information and System Security* (5:4), pp. 438 - 457.

Grambsch, P. M., and Therneau, T. M. 1994. "Proportional Hazards Tests and Diagnostics Based on Weighted Residuals," *Biometrics* (81), pp. 515-526.

Heckman, J. J. 1979. "Sample Selection Bias as a Specification Error," *Econometrica* (47:1), pp. 153-161.

iDefense Labs. 2010. "iDefense Labs Legal Notices, Terms & Conditions and Policies" (http://labs.idefense.com/legal.php).

Johansson, J. K. 1979. "Advertising and the S-Curve: A New Approach," *Journal of Marketing Research* (16:3), pp. 346-354.

Jones, J. R. 2007. "Estimating Software Vulnerabilities," *IEEE Security and Privacy* (5:4), pp. 28-32.

Kannan, K., and Telang, R. 2005. "Market for Software Vulnerabilities? Think Again," *Management Science* (51:5), pp. 726-740.

Kemmerer, R. A., and Vigna, G. 2002. "Intrusion Detection: A Brief History and Overview " *IEEE Computer* (35:4), pp. 27-30.

Kotenko, I., and Stepashkin, M. 2006. "Attack Graph Based Evaluation of Network Security," in *Communications and Multimedia Security*, H. Leitold and E. Markatos (eds.), Berlin: Springer, pp. 216-227.

Kshetri, N. 2005. "Pattern of Global Cyber War and Crime: A Conceptual Framework," *Journal of International Management* (11), pp. 541-562.

Lemos, R. 2010. "NSS Labs' Exploit Hub, a Marketplace Where Coders Can Sell Attacks on Specific Vulnerabilities, Could Help with Enterprise Security," Security Dark Reading (www.darkreading.com)

Lewis, J. B., and Linzer, D. A. 2005. "Estimating Regression Models in Which the Dependent Variable Is Based on Estimates," *Political Analysis* (13:4), 2005, pp. 345-364.

Li, P., and Rao, H. R. 2007. "An Examination of Private Intermediaries' Roles in Software Vulnerabilities Disclosure," *Information Systems Frontiers* (9:5), pp. 531-539.

Lieberman, M. B. 1987. "The Learning Curve, Diffusion, and Competitive Strategy," *Strategic Management Journal*), pp. 441-452.

Lohmeyer, D. F., McCrory, J., and Pogreb, S. 2002. "Managing Information Security," *McKinsey Quarterly* (Special Edition 2), pp. 12-16.

Mahajan, V. 1985. *Models of Innovation Diffusion*, Newbury Park, CA: Sage Publications.

Mahmood, M. A., Siponen, M., Straub, D., Rao, H. R., and Raghu, T. S. 2010. "Moving Toward Black Hat Research in Information Systems Security: An Editorial Introduction to the Special Issue," *MIS Quarterly* (34:3), pp. 431-433.

Mell, P., and Romanosky, S. 2007. "A Complete Guide to the Common Vulnerability Scoring System Version 2.0," Forum of Incident Response and Security Teams (http://www.first.org/cvss/cvss-guide.html).

Mell, P., Scarfone, K., and Romanosky, S. 2006. "Common Vulnerability Scoring System," *IEEE Security and Privacy* (4:6), pp. 85-89.

Mullin, W. P. 2001. "Will Gun Buyback Programs Increase the Quantity of Guns?," *International Review of Law & Economics* (21:1), pp. 87-102.

NVD. 2008. "National Vulnerability Database," National Institute of Standards and Technology (http://nvd.nist. gov/).

Ozment, A. 2004. "Bug Auctions: Vulnerability Markets Reconsidered," paper presented at the Workshop on the Economics of Information Security, Minneapolis, MN, May 13-14.

Plumper, T., and Troeger, V. E. 2007. "Efficient Estimation of Time-Invariant and Rarely Changing Variables in Finite Sample Panel Analyses with Unit Fixed Effects," *Political Analysis* (15:2), pp. 124-139.

Radianti, J., and Gonzalez, J. J. 2007. "Understanding Hidden Information Security Threats: The Vulnerability Black Market," in *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, Los Alamitos: IEEE Computer Society Press.

Ransbotham, S., and Mitra, S. 2009. "Choice and Chance: A Conceptual Model of Paths to Information Security Compromise," *Information Systems Research* (20:1), pp. 121-139.

Rogers, E. M. 2003. *Diffusion of Innovations* (5th ed.), New York: The Free Press.

Sandhu, R., and Samarati, P. 1996. "Authentication, Access Control, and Audit," *ACM Computing Surveys* (28:1), pp. 241-243.

Saxonhouse, G. R. 1976. "Estimated Parameters as Dependent Variables," *American Economic Review* (66:1), pp. 178-183.

Schechter, S. 2004. "Computer Security, Strength and Risk: A Quantitative Approach," unpublished Ph.D. Thesis, Harvard University (http://www.eecs.harvard.edu/~stuart/papers/thesis.pdf).

Schultz, E. 2004. "Sarbanes-Oxley: A Huge Boon to Information Security in the US," *Computers & Security* (23:5), pp. 353-354.

Siponen, M. 2005. "Analysis of Modern IS Security Development Approaches: Towards the Next Generation of Social and Adaptable ISS Methods" *Information and Organization* (15), pp. 339-375.

Straub, D., and Welke, R. 1998. "Coping with Systems Risk: Security Planning Models for Management Decision Making," *MIS Quarterly* (22:4), pp. 441-469.

Sutherland, E. 1947. *Principles of Criminology*, Philadelphia: Lippencot.

Telang, R., and Wattal, S. 2007. "An Empirical Analysis of the Impact of Software Vulnerability Announcements on Firm Stock Price," *IEEE Transactions on Software Engineering* (33:8), pp. 544-557.

Thomlinson, M. 2010. "Announcing Coordinated Vulnerability Disclosure," Microsoft Security Response Center, Microsoft Technet Blogs (http://blogs.technet.com/b/msrc/)

Tipping Point. 2010. "Zero Day Initiative" (http://www.zerodayinitiative.com/).

Unsecurity Research. 2010. "Vulnerability Marketplace Survey" (http://unsecurityresearch.com/).

Van den Bulte, C. 2000. "New Product Diffusion Acceleration: Measurement and Analysis," *Marketing Science* (19:4), pp. 366-380.

Van den Bulte, C., and Joshi, Y. V. 2007. "New Product Diffusion with Influentials and Imitators," *Marketing Science* (26:3), pp. 400-421.

Van den Bulte, C., and Stremersch, S. 2004. "Social Contagion and Income Heterogeneity in New Product Diffusion: A Meta-Analytic Test," *Marketing Science* (41:4), pp. 530-544.

About the Authors

Sam Ransbotham is an assistant professor at the Carroll School of Management at Boston College. His current research focuses on IT security and technology strategy. He received his Ph.D. from Georgia Tech. His research has appeared or is forthcoming in *Information Systems Research, Management Science*, and *INFORMS Journal on Computing*. He has also published several research articles in refereed conference proceedings including the International Conference on Information Systems, Workshop on Information Technology and Systems, and the Workshop on the Economics of Information Security.

Sabyasachi Mitra is the William H.Anderson II Associate Professor of Information Technology Management and the faculty director for the Executive MBA program at the College of Management at Georgia Tech. His current research interests focus on IT security and business continuity, managing innovation in IT industries, and the use of IT in agricultural supply chains. His research has ap-

peared or is forthcoming in journals such as *Information Systems Research*, *Management Science, MIS Quarterly, Journal of Marketing, Journal of Operations Management*, *INFORMS Journal on Computing*, *Journal of MIS*, and others. He received his Bachelor of Technology degree from the Indian Institute of Technology and his Ph.D. in Business Administration from the University of Iowa.

Jon Ramsey is the Chief Technology Officer of SecureWorks, Inc., a managed security service provider. Jon has 12 years of hands-on experience at many levels: system administrator, software engineer, analyst, security penetration specialist and senior engineer. Prior to joining SecureWorks, Jon worked for the Computer Emergency Response Team (CERT), Siemens, and the University of Pittsburgh.