

Durham Research Online

Deposited in DRO:

08 February 2021

Version of attached file:

Draft Version

Peer-review status of attached file:

Unknown

Citation for published item:

Allodi, Luca and Massacci, Fabio and Williams, Julian 'The Work-Averse Cyber Attacker Model: Theory and Evidence From Two Million Attack Signatures.', Working Paper. UNSPECIFIED.

Further information on publisher's website:

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2862299

Publisher's copyright statement:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

The Work-Averse Cyber Attacker Model: Theory and Evidence From Two Million Attack Signatures

Allodi, Luca
Technical University of Eindhoven
l.allodi@tue.nl

Massacci, Fabio
University of Trento
fabio.massacci@unitn.it

Williams, Julian
Durham University Business School
julian.williams@durham.ac.uk

Abstract

A common presumption is that the typical cyber attacker is assumed to exploit all possible vulnerabilities with almost equal likelihood. That is, the probability of an attack on a given vulnerability is at maximum entropy, one cannot importance sample which vulnerability will be exploited first, hence decision making is purely a function of criticality. In this paper we present, and empirically validate, a novel and more realistic attacker model. The intuition of our model is that a mass attacker will optimally choose whether to act and weaponize a new vulnerability, or keep using existing toolkits if there are enough vulnerable users. The model predicts that mass attackers may i) exploit only one vulnerability per software version, ii) include only vulnerabilities with low attack complexity, and iii) be slow at introducing new vulnerabilities into their arsenal. We empirically test these predictions by analysing data collected on attacks against more than one million real systems by Symantec’s WINE platform. Our analysis shows that mass attackers’ fixed costs are indeed significant. Significant efficiency gains can be made by individuals and organizations by accounting for this effect.

Keywords: Cyber Security; Dynamic Programming; Malware Production; Risk Management.

1 Introduction

Security vulnerabilities in an information system allow cyber attackers to exploit, infiltrate and exfiltrate information for financial and/or political gain. Whilst a great deal of prior research has focused on the security investment decision making process for security vendors and targets, the choices of attackers are less well understood. This paper is the first attempt to define a cyber attacker revenue function with costly effort from first principles and then to empirically fit this function to a large data set.

Our main results challenge the notion of the all powerful attackers willing to exploit a broad range of security vulnerabilities and provide important guidance to policy makers and potential targets on how to utilize finite resources in the presence of cyber threats.¹

¹An early security reference on this conceit can be found in (Dolev and Yao, 1983a, page 199) where a protocol should be secured “against arbitrary behavior of the saboteur”. The Dolev-Yao model is a quasi “micro-economic” security model: given a population with several (interacting) agents, a powerful attacker will exploit any weaknesses,

A natural starting point when attempting to evaluate the decision making of attackers is to look at the actual ‘traces’ their attacks leave on real systems (also referred as ‘attacks in the wild’): each attempt to attack a system using a vulnerability and an exploit mechanism generates a specific *attack signature*, which may be recorded by software security vendors. Dumitras and Shou (2011) and Bilge and Dumitras (2012) provide a summary of signature identification and recording, whilst Allodi and Massacci (2014) shows how they can be linked to vulnerabilities that attackers seek to exploit. By observing the frequency with which these attack signatures are triggered, it is possible to estimate (within some level of approximation) the rate of arrival of new attacks. Evidence from past empirical studies suggests a different behavior depending on fraud type; for example, Murdoch et al. (2010) shows that attackers focussing on chip and pin credit cards, which require physical access, are very proactive and rapidly update their menu of exploits; for web users, Allodi and Massacci (2014) and Nayak et al. (2014) indicate that the actual risk of attacks in the wild is limited to hundred vulnerabilities out of the fifty thousand reported in vulnerability databases. Mitra and Ransbotham (2015) confirm these findings by showing that even (un)timely disclosures do not correlate with attack volumes.

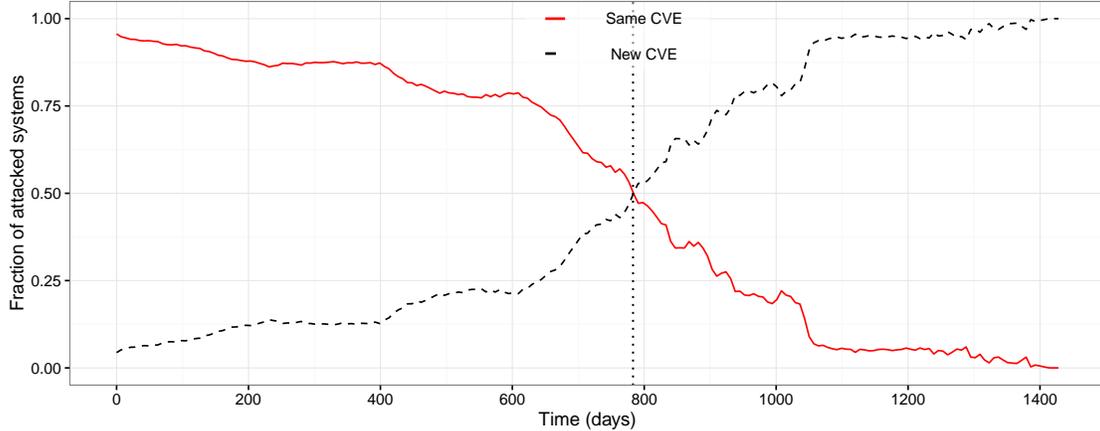
The current cannon of results provide strong prima-facie evidence that attackers do not quickly mass develop new vulnerability exploits that supplement or replace previous attack implementations. This collective behavior is at odds with the classical theoretical models of attacker behaviour Dolev and Yao (1983b): attackers can and will exploit *any* vulnerability. We must therefore conclude that attackers are rational, that the effort required to produce an exploit and hence deployable malware is costly, and that they will respond to incentives in a way that is consistent with classical models of behaviour (Laffont and Martimort, 2009). This work addresses this observation and impacts the development of risk models for cyber-attacks by identifying empirical and theoretical aspects of attack and malware production.²

Figure 1 shows the fractions of systems receiving attacks recorded by Symantec, a large security vendor, for two different cases: the red line plots the fraction of systems receiving two attacks at two different times that target the same software vulnerability (CVE). The abscissa values represent the time, in days, between attacks, hence we would expect that the red line would decrease (which it does) from near unity to zero. The black dashed line represents the dual case: the same system and the same software are attacked but the attacker uses a new vulnerability, different from the original attack. The data suggests that it takes more than two years before the number of attacks using new vulnerabilities exceeds the number of attacks using the same vulnerabilities, and about 3-4 years before a complete refresh occurs.

Our methodological contribution is twofold. First, we specify a novel theoretical model of the dynamic decision making process of the attacker that is based on Stokey’s logic of inaction, see Stokey (2008). We model the timing of effort by the attacker as a dynamic programming

and security is violated if they can compromise some honest agent. So, *every* agent must be secured by mitigating *all* vulnerabilities. An alternative formulation is that the likelihood that a given vulnerability is exploited should be at maximum entropy and hence the only dominating factor will be the criticality of that vulnerability.

²Whilst challenging the maximum entropy notion of attacks (if it can, it will) may appear to the informed reader as attacking a straw man, the underpinning ideas of Dolev-Yao persist through to present security practice, see for instance the comments in Schneier (2008) that cover similar ground. Variants of the all-powerful attackers are proposed (e.g. honest-but-curious, game-based provable security models) but they only changed the power and speed of attacks not the will: if there is weakness that the attacker can find and exploit, they will do it. Papers analyzing web vulnerabilities Stock et al. (2013); Nikiforakis et al. (2014) report statistics on the persistence of these vulnerabilities on internet sites as evidence for this all powerful effect and broad coverage of security vulnerabilities. A better understanding of the rationale behind attack might lead to more informed decision.



Fraction of systems receiving the same attack repeatedly in time (red, solid) compared to those receiving a second attack against a different vulnerability (black, dashed). The vertical line indicates number of days after the first attacks where it is more likely to receive an attack against a new vulnerability rather than against an old one.

Figure 1: Distribution of time between of subsequent attacks with similar signatures.

problem, which we initially solve in its generality. For the purpose of empirical analysis we then restrict it to the case of an attacker focusing on the ‘next’ update.

Our main prediction is that the complexity of implementation of a vulnerability endogenously interacts with the update time and attackers will predominantly flock to the low hanging fruit of vulnerabilities with low complexity and high impact. Second, we capture such relations by deriving, directly from the theoretical model, a corresponding robust parametric regression model of equilibrium update times and hence reverse engineer part of the malware production function, a first in this literature. To ensure that the causal relations predicted in the theoretical model are captured in our empirical study, we control for several factors related to the characteristics of the user and their system for each recorded pair of attacks in the data (e.g. user geographical locations). This work is the first to explicitly consider an attacker with fixed-costs, and to validate the theoretical predictions with an empirical analysis derived directly from the analytical model.

The next section (§2) provides a brief review of the current research on attacker technologies and production functions. Then we provide a dynamic model of attacking effort in continuous time with discrete update intervals (§3), this model is then used to compute a linear equilibrium between updates of attackers technologies and attack intensity (measured by machines attacked). We utilize the specific insights from this general model to derive several hypothesis amenable to empirical verification. The WINE dataset spanning two millions of attack signatures recorded in the wild by Symantec is introduced (§4) and extract a number of regression models from the prime principles (§5.1) by a mathematical transformation of the general model introduced in Section 3. This approach is used to derive additional empirical hypothesis for the regression variables (Summarized in Table 8). We finally discuss the results of the empirical analysis (§5) and conclude the paper by outlining implications for theory and practice (§6).

2 Background

The economic decision making and the organization of information system security has been explored from a target perspective quite thoroughly in recent years. Anderson (2008) provides a good summary on the early extant literature.³ Commonly, the presence of vulnerabilities has been considered as the starting point for the analysis, as it potentially allow an attacker to take (total or partial) control of a system and subvert its normal operation for personal gain.

Asghari et al. (2013); Van Eeten and Bauer (2008); Van Eeten et al. (2010) address the economic incentives (perverse or well aligned) in addressing security threats. Market based responses, such as bug bounty programs, are discussed in Miller (2007) and Frei et al. (2010) whilst Karthik Kannan (2005) and Finifter et al. (2013) also address the issue of vulnerability markets and the value of vulnerabilities to developers, organizations that may potentially be targeted, and indeed attackers. Policing aspect on the disclosure of vulnerabilities, and their effect on overall security and/or availability of attacks to attackers has been investigated by multiple studies (Mitra and Ransbotham (2015); Ransbotham et al. (2012); Arora et al. (2008)). For example, Ransbotham et al. (2012) finds that the introduction of incentives for vulnerability disclosure tend to decrease exploitation attempts and increase overall security. However, Mitra and Ransbotham (2015) finds early disclosure has no impact on attack volume (number of recorded attacks) and there is only some correlation between early disclosure and the ‘time-to-arrival’ of exploits.

Defender strategies in terms of patching times have been investigated both empirically (Arora et al. (2004); Okhravi and Nicol (2008)) and theoretically (Cavusoglu et al. (2006); Serra et al. (2015b)), assuming specific threat models and attack production functions. Kannan et al. (2016) considers how software vendors may maximize profit by exploiting the (endogenously or exogenously-defined) attacker behavior; August et al. (2014) evaluates the externalities caused by outsourced security management, and find that under some circumstances diversification, as opposed to patching, may yield lower costs. Dey et al. (2015) finds that ‘single-metric’ patching policies may be ineffective, and that additional parameters and metrics may be needed. Similarly, Lee et al. (2016) find that security best practices do not necessarily lead to more robust firm security, and that different managerial settings may affect the relation between liability and security. These findings are also well-reflected in empirical studies (Bozorgi et al. (2010); Allodi and Massacci (2014); Nayak et al. (2014)), but a theoretical explanation capable of predicting the appearance of new exploits is still amiss.

Presence of *vulnerability* if oftentimes largely undifferentiated in the strictly empirical literature from presence of *exploit* due to the lack of an explicitly or implicit attacker decision model underlying the analysis (Naaliel et al. (2014); Bozorgi et al. (2010)). The theoretical aspects behind the data generation process are explored in a growing number of studies in the recent literature. These studies consider exploitation as a costly factor that can significantly affect an attacker’s decision to launch an attack or wait for better conditions to emerge (Ransbotham and Mitra (2009); van Dijk et al. (2013); Smeets (2018)). For example, van Dijk et al. (2013) models the attack/defense decision process with varying attacker strategies ranging from fixed attack updates to adaptive strategies based on the defender’s decisions. Similarly, Smeets (2018) models attacker decisions based on the expectations of the attack’s *persistence* and *stealthiness* to defender detection and remediation capabilities. These applications generally consider the attacker to be potentially capa-

³Policy, cost sharing and incentives have also been comprehensively explored from a target perspective in August and Tunca (2006, 2008, 2011) and Arora et al. (2004, 2008).

ble of adopting any strategy with different degrees of probability, depending on the conditions of the game (Manshaei et al. (2013)). This holds particularly relevant for targeted attacks scenarios whereby attackers may accurately measure the state of a specific target and decide whether to employ or engineer a (new) attack based on their belief that this will maximize some variable of value (e.g. extended control over a resource as in van Dijk et al. (2013)). Differently, attacks against large pools of ‘similar’ targets (e.g. by geographical distribution, or system configuration) adapt to the state of the *population* of potential targets (as opposed to one specific instance), for which attack technologies developed ‘ad-hoc’ are not always viable (Ransbotham and Mitra (2009)).

An example from recent studies illustrates the development of an underground economy that fuels the technological development required to deliver attacks at scale, whereby changing vulnerability conditions are reflected at market level, and not at level of the single attacker (Grier et al. (2012)). This difference has been acknowledged in the modeling literature (e.g. Laszka et al. (2013)), but an explicit characterization of the attacker decision model underlying this effect remains unexplored, despite this type of attack carrying a disproportionately large fraction of the risk suffered by the final users (Provos et al. (2008); Allodi (2015); Nayak et al. (2014)). We fill this gap by providing a theoretical and empirical link between timing of attacks at scale and appearance of new exploits targeting the *mass* of Internet users.

A precise connotation of exploitation timings remains hard to develop as empirical data identifying attacker decisions remains exceedingly rare, with few exceptions focusing on specific case scenarios (Ransbotham and Mitra (2009); Mitra and Ransbotham (2015)). A few studies did propose to evaluate market and economic effects on attacked systems (Anderson and Moore (2006); Ransbotham et al. (2012); Finifter et al. (2013)), but did not tackle the underlying problem of attack production at scale. Similarly, recent developments modeled and quantified the risk relation between attack generation and defenses (Allodi and Massacci (2017); Bilge et al. (2017)), but lack of a prediction mechanism for the rate of arrival of new attacks (as opposed to new vulnerabilities).

Our study fills this gap by providing a long-missing link between theoretical models of the attacker (typically assuming all vulnerabilities will lead to an attack with a certain probability), and the contrasting empirical observation that most attacks are similar in nature and focused on few target vectors only Allodi (2015); Nayak et al. (2014).

The consensus in the security literature appears to have settled on the view that the severity of the vulnerability in terms of the criticality of the software for the functioning of the business, as measured by a series of metrics, should be used as a direct analogue of risk. For a broad summary of the metrics and risk assessments involved in this domain see Mellado et al. (2010) (technology based metrics); Sheyner et al. (2002); Wang et al. (2008); Manadhata and Wing (2011) (attack graphs and surfaces); and Naaliel et al. (2014); Wang et al. (2009); Bozorgi et al. (2010); Quinn et al. (2010) (severity measure metrics and indexation). From an economic perspective Ransbotham and Mitra (2009) studies the incentives that influence the diffusion of vulnerabilities and hence the opportunities of the attacker to attack a target’s systems (software and infrastructure). The standard de facto metric for the assessment of vulnerability severity is the Common Vulnerability Scoring System, or CVSS⁴ in short CVSS-SIG (2015). Hence, the view is that any open vulnerability will eventually be exploited by some form of malware and the target organization will then be subject to a cyber attack.

⁴The CVSS score provides a standardized framework to evaluate vulnerability severity over several metrics, and is widely reported in public vulnerability databases such as the National Vulnerability Database NIST (2015) maintained by the National Institute of Standards in Technology (NIST).

Recent studies in the academic literature have challenged the automatic transfer of the technical assessment of the ‘exploitability’ of a vulnerability into actual attacks against end users. Bozorgi et al. (2010) and Allodi and Massacci (2014) have empirically demonstrated on different samples a substantial lack of correlation between the observed attack signatures in the wild and the CVSS type severity metrics. The current trend in industry is to use these values as proxies demanding immediate action (see Beattie et al. (2002) for operating system security, PCI-DSS (2010) for credit card systems and Quinn et al. (2010) for US Federal rules).

In particular, prior work suggests that only a small subset of vulnerabilities are actually exploited in the wild (Allodi and Massacci, 2014), and that none of the CVSS measures of severity of impact predict the viability of the vulnerability as a candidate for an implemented exploit (Bozorgi et al., 2010). In particular, Bozorgi et al. (2010) argue that besides the ‘technical’ measure of a vulnerabilities, one should also consider the value or cost of a vulnerability exploit and the ease and cost with which the exploit can be developed and then deployed.

On a similar line, Herley (2013) posits the idea that for a (rational) attacker not all attack types make sensible avenues for investment. This is supported by empirical evidence showing that attack tools actively used by attackers embed only ten to twelve exploits each on the maximum (Kotov and Massacci, 2013; Grier et al., 2012), and that the vast majority of attacks recorded in the wild are driven by only a small fraction of known vulnerabilities (Nayak et al., 2014; Allodi, 2015). It is clear that for the attacker some reward must be forthcoming, as the level of costly effort required to implement and deliver the attack observed in the wild is far from negligible. For example, Grier et al. (2012) uncovers the presence of an underground market where vulnerability exploits are rented to attackers (‘exploitation-as-a-service’) as a form of revenue for exploit writers. Liu et al. (2005) suggest that attacker incentives should be considered when thinking about defensive strategies: increasing attack costs or decreasing revenue may be effective in deterring the development and deployment of an attack. For example, Chen et al. (2011) suggests to ‘diversify’ system configurations so that the fraction of attackable system by a single exploit diminishes, hence lowering the return per attack.

The effort needed to engineer an attack can be generally characterized along the three classic phases from Jonsson and Olovsson (1997): *reconnaissance* (where the attacker identifies potential victims), *deployment* (the engineering phase), *refinement* (when the attack is updated). The first phase is covered by the works of Wang et al. (2008); Howard et al. (2005); Nayak et al. (2014), where the attacker investigates the potential pool of targets affected by a specific vulnerability by evaluating the attack surface of a system, or the ‘popularity’ of a certain vulnerable software. The engineering aspects of an exploit can be understood by investigating the technical effort required to design one (see for example Schwartz et al. (2011) and Carlini and Wagner (2014) for an overview of recent exploitation techniques). However, the degree of re-invention needed to update an exploit and the anticipated time from phase two to phase three remain largely un-investigated (Yeo et al., 2014; Serra et al., 2015a; Bilge and Dumitras, 2012, provide some useful results in this direction).

Our model aggregates deployment and reconnaissance costs into a single measure, whereas we explicitly model the expected time to develop exploits and subsequent updating times to include new vulnerabilities in the arsenal.

3 A Dynamic Attacker Model with Costly Effort

We consider a continuous time setting, such that $0 < t < \infty$, where an attacker will be choosing an optimal update sequence $0 < T_1 < T_2 < \dots T_i \dots T_\infty$ for weaponizing new vulnerabilities v_1, \dots, v_n . The attacker's technology has a "combination" of exploit technology that undergoes periodic updates. Each combination targets a specific mix of vulnerabilities and we presume that attackers can make costly investments to develop the capability of their technology.

The attacker starts activity at time $t = 0$ by initially identifying a set of vulnerabilities $V \subset \mathcal{V}$ from a large universe \mathcal{V} affecting a large number of target systems N . A fraction θ_V of the N systems is affected by V and would be compromised by an exploit in absence of security countermeasures. Targets are assumed to deploy patches and/or update systems, whilst security products update their signatures (e.g. antiviruses, firewalls, intrusion prevention systems). Hence, the number of infected systems available for exploit will decay with time.

We can represent vulnerability patches and security signatures as arriving on users' systems following two independent exponential decay processes governed respectively by the rates λ_p and λ_{sig} . The effect of λ_p and λ_{sig} on attacks has been previously discussed by Arora et al. (2004, 2008); Chen et al. (2011), whilst a discussion on their relative magnitude is provided in Nappa et al. (2015). Assuming that the arrival of patches and antivirus signatures are independent processes, and rolling them up into a single factor $\lambda = f(\lambda_p, \lambda_{sig})$, the number of systems impacted by vulnerabilities in V at time t is

$$N_V(t) = N\theta_V e^{-\lambda t}. \quad (1)$$

For the given set of vulnerabilities V targeted by their technologies combination the attacker will pay an upfront cost $C(V|\emptyset)$ and has an instantaneous stochastic profit function of

$$\Pi_V(t) = [r(t, N_V(t), V) - c(t, V)]e^{-\delta t}, \quad (2)$$

where $r(t, N_V, V)$ is a stochastic revenue component,⁵ whilst $c(t, V)$ is the variable costs of maintaining the attack,⁶ subject to a discount rate of δ . We do not make any assumption on the revenues from successful attacks. They could be kudos in hackers fora (Ooi et al., 2012) or revenues from trading victim's assets in black markets (Grier et al., 2012; Allodi et al., 2015).

At some point, the attacker might decide to perform a refresh of its attacking capabilities by introducing a new vulnerability and engineering its exploit by incurring an upfront cost of $C(v|V)$. This additional vulnerability will produce a possibly larger revenue $r(t, N_{V \cup \{v\}}(t), V \cup \{v\})$ at an increased marginal cost $c(t, V \cup \{v\})$. As the cost of engineering an exploit is large with respect to maintenance ($C(v|V) \gg c(t, V \cup \{v\})$) and neither successful infection (Allodi et al., 2013), nor revenues are guaranteed (Herley and Florencio, 2010; Rao and Reiley, 2012; Allodi et al., 2015), the attacker is essentially facing a problem of deciding action vs inaction in presence of fixed initial costs as described by Stokey (2008) and one's best strategy is to deploy the new exploit only when the old vulnerabilities no longer guarantee a suitable expected profit.

This decision problem is then repeated over time for n newly discovered vulnerabilities, and n refresh times denoted by T_i .

⁵This component accounts for the probability of establishing contact with vulnerable system (Franklin et al., 2007), the probability of a successful infection given a contact (Kotov and Massacci, 2013; Allodi et al., 2013), and the monetization of the infected system (Kanich et al., 2008; Zhuge et al., 2009; Rao and Reiley, 2012).

⁶For example, attackers may need to obfuscate the attack payload to avoid detection (Grier et al., 2012), or renew the domain names that the malware contacts to prevent domain blacklisting (Stone-Gross et al., 2009).

We define by $C_0 = C(V|\emptyset)$ the initial development cost and by $C_{i+1} \equiv C(v_{i+1}|V \cup \{v_1 \dots v_i\})$ the cost of developing the new exploits, given the initial set V and the additional vulnerabilities $v_1 \dots v_i$. We denote by $N_i(t) \equiv N_{V \cup \{v_1, \dots, v_i\}}(t)$ the number of systems affected by adding the new vulnerability at time t . Similarly, we define $r_i(t)$ and $c_i(t)$ as the respective revenue and marginal cost of the vulnerability set $V \cup \{v_1, \dots, v_i\}$. Then the attacker faces the following stochastic programming problem for $n \rightarrow \infty$

$$\{T_1^*, \dots, T_n^*\} = \arg \max_{\{T_1, \dots, T_n\}} \sum_{i=0}^n -C_i e^{-\delta T_i} + \int_{T_i}^{T_{i+1}} (r_i(t, N_i(t)) - c_i(t)) e^{-\delta t} d\omega. \quad (3)$$

The action times $T_0 = 0$, $T_{i+1} > T_i$, and T_{n+1} is such that $r_n(T_{n+1}, N_n(T_{n+1})) = c_n(T_{n+1})$. Since the maintenance of malware, for example through ‘packing’ and obfuscation (i.e. techniques that change the aspect of malware in memory to avoid detection), is minimal and does not depend on the particular vulnerability, see (Brand et al., 2010, § 3) for a review of the various techniques, we have that $c_i(t) \rightarrow 0$ and therefore also $T_{n+1} \rightarrow \infty$. This problem can be solved with the techniques discussed in Stokey (2008), Birge (2010), Rincón-Zapatero and Rodríguez-Palmero (2003) and further developed in Birge and Louveaux (2011) either analytically or numerically by simulation. Some further mild assumptions result in solutions with a clean set of predictions that motivate and place in context our empirical work in the standard Markovian set-up needed to identify specific effects.

By imposing an instantaneous, history-less payoff (adapted process), with a risk neutral preference, expected pay-off and expected utility for a given set of ordered action times $\{T_1, \dots, T_n\}$ coincide. Risk preferences are then encapsulated in the discount factor, a common assumption in dynamic programming literature (see the discussion on model choice in (Stokey, 2008, Ch. 1)).

Hence, the simplest approach is to presume risk neutrality (under the discount factor δ) and solve in expectations as a non-stochastic Hamilton–Jacobi–Bellman type problem along the standard principles of optimal decision making (Birge and Louveaux, 2011, Ch. 4). Under the assumption of stationary revenues, we define r as the average revenue across all systems. The expected payoff from deployed malware at time t (where $t \geq T$ is the amount of time since the attacker updated the menu of vulnerabilities by engineering new exploits at time T) is then approximated by definition as follows:

$$\mathbb{E}[r(t, N_{V \cup \{v\}}(t))] \stackrel{def}{=} rN \left(\theta_V e^{-\lambda t} + (\theta_{V \cup \{v\}} - \theta_V) e^{-\lambda(t-T)} \right), \quad (4)$$

The first addendum caters for the systems vulnerable to the set V of exploited vulnerabilities that have been already partly patched, whilst the second addendum accounts for the new, different systems that can be exploited by adding v to the pool. For the latter systems, the un-patched fraction restarts from one at time T .

Solving Eq. (3) in expectations we can replace the stochastic integral over $d\omega$ with a traditional Riemann integral over dt and evaluate the approximation in expectations. By solving the above integral and imposing the usual f.o. condition we obtain a decomposition for the optimal profit for the attacker.

$$\{T_1^*, \dots, T_n^*\} = \arg \max_{\{T_1, \dots, T_n\}} \sum_{i=0}^n (\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i} \quad (5)$$

$$\Pi(T_{i+1}, T_i) = \frac{rN}{\lambda + \delta} \left(\theta_i - \theta_{i-1} + \theta_{i-1} e^{-\lambda T_i} \right) \left(1 - e^{-(\lambda + \delta)(T_{i+1} - T_i)} \right) \quad (6)$$

where we abbreviate $\theta_{-1} \equiv 0$, $\theta_0 \equiv \theta_V$, and $\theta_i \equiv \theta_{V \cup \{v_1 \dots v_i\}}$.

Proposition 1. *The optimal times to weaponize and deploy new exploits for attackers aware of initial fixed costs of exploit development results by solving the equations below for $i = 1, \dots, n$.*

$$\frac{\partial \Pi(T_i, T_{i-1})}{\partial T_i} e^{-\delta T_{i-1}} - \delta(\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i} + \frac{\partial \Pi(T_{i+1}, T_i)}{\partial T_i} e^{-\delta T_i} = 0 \quad (7)$$

subject to $T_0 = 0$, $T_{n+1} = \infty$, $\delta, \lambda > 0$.

Proof of Proposition 1 is given in Appendix 7.1. □

Unrestricted dynamic programming problems such as that described in Proposition 1 do not generally have analytic solutions for all parameter configurations. They can be log-solved either in numerical format, or by computer algebra as a system of n equations by setting $x_i = e^{-\lambda T_i}$, $y_i = e^{-\delta T_i}$ and then adding the n equations $\delta \log x_i = \lambda \log y_i$. However, we can obtain a closed form solution in the case of a admit a myopic attacker who only considers a single individual decision $n = 1$. This approximates the case when the true dynamic programming problem results in $T_1^* > 0$ and $T_2^* \rightarrow \infty$. For an overview of the appropriate domains for the use of this type of simplification see DeGroot (2005).

Corollary 1. *A myopic attacker, who anticipates an adapted revenue process from deployed exploits subject to a decreased effectiveness due to patching and anti-virus updates with a negligible cost of maintenance for each exploit, will postpone indefinitely the choice of weaponizing a vulnerability v if the ratio between the cost of developing the exploit and the maximal marginal expected revenue is larger than the discounted increase in the fraction of exploited vulnerabilities, namely $C(v|V)/rN > \delta/(\lambda + \delta)(\theta_{V \cup \{v\}} - \theta_V)$. The attacker would be indifferent to the time at which deploy the exploit only when the above relation holds at equality.*

Proof of Corollary 1 is given in Appendix 7.2. □

Whilst our focus is on realistic update times for the creation of new exploits, note that the ‘all-powerful’ attacker is still admitted as a particular case when the attacker cost function $C(v|V)$ for weaponizing a new vulnerability collapses to zero. In this case, Corollary 1 predicts that the attacker could essentially deploy the new exploit at an arbitrary time $[0, +\infty]$ even if the new exploit would not yield extract impact.

If the vulnerability v affects a software version for which there is already a vulnerability in V , the fraction of systems available for exploit will be unchanged ($\theta_{V \cup \{v\}} = \theta_V$). Hence, the cost has to be essentially close to zero ($C(v|V) \rightarrow 0$) for the refresh to be worth. In the empirical analysis section §(5) we report a particular case where this phenomenon occurs in our dataset. Based on these considerations we can now state our first empirical prediction.

Hypothesis 1. Given Corollary 1, a work-averse attacker will overwhelmingly use only one reliable exploit per software version.

Engineering a reliable vulnerability exploit requires the attacker to gather and process technical vulnerability information (Erickson, 2008).

This technical ‘exploit complexity’ is captured by the **Attack Complexity** metric provided in the CVSS. When two vulnerabilities cover essentially the same population ($\theta_{V \cup \{v\}} - \theta_V \approx \epsilon$) a lower cost would make it more appealing for an attacker to refresh their arsenal as this would make it

easier to reach the condition $(C(v|V)/rN \approx \delta/(\lambda + \delta)(\theta_{V \cup \{v\}} - \theta_V) \approx \epsilon)$ when the attacker would consider deploying an exploit to have a positive marginal benefit.

Hypothesis 2. Corollary 1 also implies that a work-averse attacker will preferably deploy low-complexity exploits for software with the same type of popularity.

Corollary 1 describes a result in the limit and in presence of a continuous profit function. Indeed according to Eq. (4) the attacker expects to make a marginal profit per unit of time equal to $rNf(t)$ where $\lim_{t \rightarrow \infty} f(t) \rightarrow 0$ and as a result $\partial\Pi(T_{i+1}, T_i)/\partial T_{i+1}$ is a monotone decreasing function and $\partial\Pi(T_{i+1}, T_i)/\partial T_{i+1} \rightarrow 0$ for $T_{i+1}b \rightarrow \infty$. In practice, the profit expectations of the attacker are discrete: as the marginal profit drops below r , it is below the expect marginal profit per unit of compromised computers. Hence, the attacker will consider the time $T_{i+1} = T^* < \infty$ where such event happens as equivalent to the event where the marginal profit goes to zero ($T_{i+1} = \infty$) and hence assumes that the maximal revenue has been achieved and a new exploit can be deployed.

Proposition 2. *A myopic attacker, who anticipates an adapted revenue process from deployed exploits subject to a decreased effectiveness due to patching and anti-virus updates with a negligible cost of maintenance for each exploit, and expects a marginal profit at least equal to the marginal revenue for a single machine ($\partial\Pi/\partial T \geq r(0, N_V(0), V)/N_V(0)$) will renew their exploit at*

$$T^* = \frac{1}{\delta} \log \left(\frac{C(v|V)}{r} - \frac{\delta}{\lambda + \delta} (\theta_{V \cup \{v\}} - \theta_V) N \right) \quad (8)$$

under the condition that $\frac{C(v|V)}{rN} \geq \frac{1}{N} + \frac{\delta}{\lambda + \delta} (\theta_{V \cup \{v\}} - \theta_V)$.

Proof of Proposition 2 is given in Appendix 7.3. □

Assuming the cost and integral of the reward function over $[0, T_i^*]$ are measured in the same numéraire and approximately within the same order of magnitude, the model implies that the discount factor plays a leading role in determining the optimal time for the new exploit deployment, the term $\frac{1}{\delta}$ in Eq. (8). Typically the extant microeconomics literature (see Frederick et al., 2002) sets $\exp(\delta) - 1$ to vary between one and twenty percent. Hence, a lower bound on T_1^* would be $\approx [100, 400]$ when time is measured in days. This implies the following prediction:

Hypothesis 3. Given Proposition 2, the time interval after which a new exploit would economically dominate an existing exploit is large, $T_1^* > 100$ days.

4 Data Set

Our empirical dataset merges three data sources, these are:

The **National Vulnerability Database (NVD)** is the vulnerability database maintained by the US. Known and publicly disclosed vulnerabilities are published in this dataset along with descriptive information such as publication date, affected software, and a technical assessment of the vulnerability as provided by the CVSS. Vulnerabilities reported in NVD are identified by a Common Vulnerabilities and Exposures identifier (CVE-ID) that is unique for every vulnerability.

The **Symantec threat report database (SYM)** reports the list of attack signatures detected by Symantec’s products along with a description in plain English of the attack. Amongst other information, the description reports the CVE-ID exploited in the attack, if any.

Table 1: Variables included in our dataset

Variable	Description
$CVE_{1,2}$	The identifier of the previous and the current vulnerability v exploited on the user’s machine.
\mathcal{T}	The delay expressed in fraction of year between the first and the second attack.
\mathcal{N}	The number of detected attacks for the pair <i>previous attack</i> , <i>actual attack</i> .
\mathcal{U}	The number of systems attacked by the pair.
Comp1	The Complexity of the vulnerability as indicated by its CVSS assessment. Can be either High , Medium or Low as defined by CVSS(v2) Mell et al. (2007).
Imp	The Impact of the vulnerability measured over the loss in Confidentiality, Integrity and Availability of the affected information. It is computed on a scale from 0 to 10 where 10 represents maximum loss in all metrics, and 0 represents no loss. Mell et al. (2007).
Day	The date of the vulnerability publication on the National Vulnerability Database.
Sw	The name of the software affected by the vulnerability.
Ver	The last version of the affected software where the vulnerability is present.
Geo	The country where the user system is at the time of the second attack.
Hst	The profile of the user or “host”. See Table 2 for reference.
Frq	The average number of attacks received by a user per day. See Table 2.
Pk	The maximum number of attacks received by a user per day. See Table 2.

The **Worldwide Intelligence Network Environment (WINE)**, maintained by Symantec, reports *attack signatures* detected in the wild by Symantec’s products. In particular, WINE is a representative, anonymized sample of the operational data Symantec collects from users that have opted in to share telemetry data (Dumitras and Shou, 2011). WINE comprises attack data from more than one million hosts, and for each of them, we are tracking up to three years of attacks. Attacks in WINE are identified by an ID that identifies the attack signature triggered by the detected event according to Symantec’s threat database. To obtain the exploited vulnerability we match the attack signature ID in WINE with the CVE-ID reported in SYM.

The data extraction involved three phases: (1) reconstruction of WINE users’ attack history; (2) building the controls for the data; (3) merging and aggregating data from (1) and (2). Because of user privacy concerns and ethical reasons, we did not extract from the WINE dataset any potentially identifying information about its hosts. For this reason, it is useful to distinguish two types of tables: tables *computed* from WINE, namely intermediate tables with detailed information that we use to build the final dataset; and *extracted* tables, containing only aggregate information on user attacks that we use in this research. The full list of variables included in our dataset is described in Table 1.⁷

4.1 Understanding the Attack Data Records

We are interested in the new vulnerability v whose mass exploit is being attempted in the wild after an exploit for V vulnerabilities have been already engineered and attempted in the recent past. Our goal is to empirically evaluate whether this past is indeed more or less recent.

⁷Replication of our analysis can find NVD publicly available at <http://nvd.nist.gov>; SYM is available online by visiting http://www.symantec.com/security_response/landing/threats.jsp. The version of SYM and NVD used for the analysis is also available from the authors at <https://securitylab.disi.unitn.it/doku.php?id=datasets>; the full dataset computed from WINE was collected in July 2013 and is available for sharing at Symantec Research Labs (under NDA clauses for access to the WINE repository) under the reference *WINE-2012-008*. In the online Appendix B we provide a full ‘replication guide’ that interested researchers may follow to reproduce our results from similar sources by Symantec or other security vendors.

Table 2: Values of the **Hst**, **Frq**, and **Pk** control variables for WINE users.

Hst	Description	Frq/Pk	Description
STABLE	Host does not update and does not change country between attacks.	LOW	#attacks ≤ 1 .
ROAM	Host’s system is the same but it changed location.	MEDIUM	#attacks ≤ 10 .
UPGRADE	Host’s system was upgraded without moving.	HIGH	#attacks ≤ 100 .
EVOLVE	Host both upgraded the system and changed location.	VERYHIGH	#attacks ≤ 1000 .
		EXTREME	#attacks > 1000 .
		Frq	average \times day
		Pk	maximum \times day

To do so we initially (1) extract from WINE two attack signatures received by a system (host) monitored by Symantec at different moments in time, (2) associate each attack signature to the corresponding vulnerability whose exploit is attempted (Combining WINE, SYM and NVD), and (3) collect from WINE some features of the host which suffered such attacks as control variables.

We use the host’s profile in terms of countries it connects to the Internet from, whether the host moves geographically, and whether the host upgraded to a new version of the operating system because users with profiles that change in time may look different to the attacker, and may therefore be subject to different attacks and attack volumes (see Chen et al. (2011); Kotov and Massacci (2013); Grier et al. (2012); Baltazar (2011) for a discussion).

Table 2 reports the measured values for each dimension and their definition. The **Hst** variable measures whether the host changed geographical region since the first attack happened (as this may affect their likelihood of receiving an attack), and whether it update its system in the observation period. **Frq** and **Pk** measure respectively the average and maximum number of attacks received per day by the host. We use them as proxy variables measuring the ‘exposure’ of a host to attacks. Thresholds have been chosen based on attacks distribution received per day by users in WINE.

To avoid idiosyncratic noise from individual systems (as well as avoid processing sensitive data such as IP numbers) we aggregated the information across hosts having the same characteristics. Hence, each row of the dataset shows the aggregated number of attacked hosts (\mathcal{U}) and the aggregated volume of attacks (\mathcal{N}) that have targeted a vulnerability v identified by CVE_2 , such that the very same hosts have previously (\mathcal{T} days before) received an attack on another vulnerability identified by CVE_1 . Hence, an “attack-pair” describes the subsequent exploitation of two (possibly identical) vulnerabilities on a *set* of host systems with similar characteristics. This allows us to evaluate the evolution attack arrivals against comparable systems worldwide. Whilst each $\langle \text{CVE}_1, \text{CVE}_2, \mathcal{T}, \text{Geo}, \text{Hst}, \text{Frq}, \text{Pk} \rangle$ tuple is unique in the dataset, the same CVE_i may appear multiple times in the dataset, as may the same pair $\text{CVE}_1, \text{CVE}_2$.

Table 3 reports an excerpt from the dataset (we omit **Frq**, **Pk**, and all the CVSS details of CVE_1 and CVE_2 for brevity). Each row represents a pair of detected attack signature. The columns CVE_1 and CVE_2 report respectively the CVE-ID of the attacked vulnerability in v and in the novel attack against V . Column \mathcal{T} reports the time delay, measured in days, between the two attacks. Column \mathcal{N} reports the overall number of attacks detected for CVE_2 after an attack against CVE_1 ; \mathcal{U} reports the number of single systems receiving the same pair of attacks. Column **Geo** reports the country in which the second attack was recorded. Finally, **Hst** reports the type of user affected by the attack

Table 3: Summary Excerpt from our dataset.

CVE ₁	CVE ₂	\mathcal{T}	\mathcal{U}	\mathcal{N}	Geo	Hst
2003-0533	2008-4250	83	186	830	IT	Up
2003-0818	2003-0818	146	1	1	US	Rm
2003-0818	2009-4324	616	1	1	CH	Ev
2003-0818	2009-4324	70	52	55	US	Ev

For example, on the third row, one WINE system ($\mathcal{U}=1$) located in Switzerland ($\text{Geo}=\text{CH}$) suffered only once ($\mathcal{N}=1$) from an attack targeting the vulnerability $\text{CVE}_2 = \text{CVE-2009-4324}$ that was preceded by an attack targeting $\text{CVE}_1 = \text{CVE-2003-0818}$ almost two years earlier ($\mathcal{T}=616$). On the fourth row, $\mathcal{U}=52$ systems in the United States ($\text{Geo}=\text{US}$) received $\mathcal{N}=55$ times the first attack on CVE_1 on followed by the second attack on CVE_2 just two months apart ($\mathcal{T}=70$). In both cases the systems considered are of type EVOLVE, indicating that the affected systems have been upgraded and moved from some other country to the country listed in Geo during our observation period.

as defined in Table 2.

Information regarding both attacked CVEs is extracted from the NVD: for each CVE we collect publication date (Day), vulnerable software (Sw), last vulnerable version (Ver), and an assessment of the Comp1 of the vulnerability exploitation and of its Imp , provided by CVSS (v2).

As we mentioned, we associate attack signature to CVE by combining information from WINE with Symantec own database of attack signatures (SYM). However, attack signatures as reported by Symantec have varying degrees of generality, meaning that they can be triggered by attacks that targets different vulnerabilities but still follow some common pattern. For this reason, some signatures reference more than one vulnerability. In this case, we have no means to know which of the vulnerabilities was effectively targeted by the attack. Out of 1,573 different attack signatures, 112 involve more than one vulnerability; to avoid introducing counting errors on the number of attacks per CVE, we dropped these attack signatures from further consideration.

Figure 2 reports the observed distribution on a logarithmic scale. It is apparent that most WINE hosts receive only a handful of attacks per day, while few hosts are subject to ‘extreme’ levels of attack density ($> 1,000/\text{day}$). These may be hosts subject to network floods sent by the attacker in a ‘scripted’ fashion, or a large number of individual users whose actual, private IP address is behind a large proxy (e.g. by their internet service provider).

4.2 Descriptive statistics

Descriptive statistics of all variables are reported in Table 4. The upper left quadrant presents the mean and standard deviation Delay (\mathcal{T}), Volume (\mathcal{N}) and Machines (\mathcal{U}), we have 2.57 Million rows of data for each pair of attacks. The average delay is about a year, whilst the associated volume, the number of attacks for the pair, is about thirteen and a half, against an average \mathcal{U} count of eleven. However, whilst the standard deviation of \mathcal{T} is less than the mean, the \mathcal{N} and \mathcal{U} counts are substantially larger, as both are truncated at zero, the right tail is very large.

The left lower quadrant of Table 4, in three parts, presents the counts for the dummy variables generated from the categorical variables that represent the user profile of the user or host (Hst), the number of attacks received by the user per day (Frq) and the maximum number of attacks received (Pk). We can see that the EVOLVE option in Hst is the most common with 1.42 million observations from 2.57 million attack pairs. For Frq the LOW and MEDIUM categories are the most common at around one million occurrences each. Finally Pk is dominated by the medium

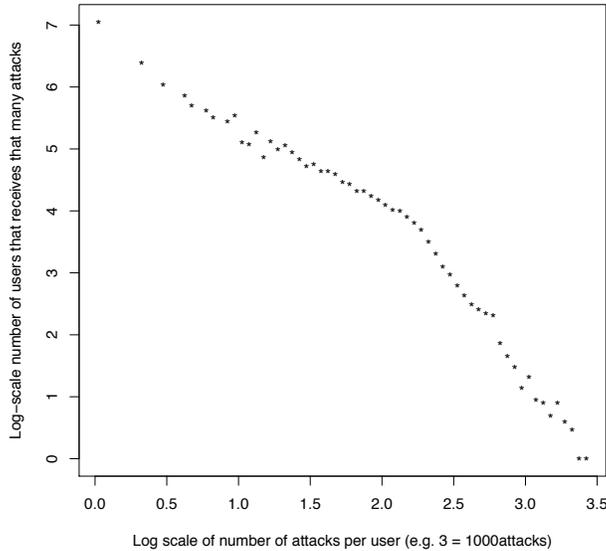


Figure 2: Distribution of attacks received per day by WINE users.

Note: Scatterplot distribution of attack frequency per user. Log-number of users is reported on the y-axis; the x-axis reports (log) frequency of attacks. There is an exponential relationship between the variables: few users receive thousands of attacks (> 1000 in the observed period), as opposed to the vast majority that only receive a few.

category with 1.57 million occurrences, about 60% of the observations in this category.

The right half of Table 4 presents the geographical split of the attack pairs. As expected Western Europe (15%) and North America (47%) represent about two thirds of the attack pairs. However, a substantial number, 29%, did not have an available geographical location. It is notable that Eastern Asia and Australia & New Zealand do represent about a tenth of the observations in the sample. Summary statistics of the collected variables are reported in Table 5.

To look into the specific properties \mathcal{N} by geography in particular, Figure 3 reports the distribution of the mean number of attacked systems per day in each geographic area of the five continents. The distributions are significantly different among and within continents with the exception of Africa, for which we observe little intra-continental variance. The highest mean arrival of attacks per day is registered in Northern America, Asia and throughout Europe with the exception of Southern Europe. The mapping of country and region is defined as in the World Bank Development Indicators.⁸

5 Empirical Analysis

The data is quite obviously unique and hence prior to conducting any correlative analysis we illustrate some scenarios that provide prima facie statistical evidence on the validity of the hypotheses identified from our theoretical model. In accordance with Hyp. 1 the attacker should prefer to (a) attack the same vulnerability multiple times rather than for only a short period of time, and (b) create a new exploit only when they want to attack a new software version.

⁸See <http://data.worldbank.org/country> for a full categorization and breakdown.

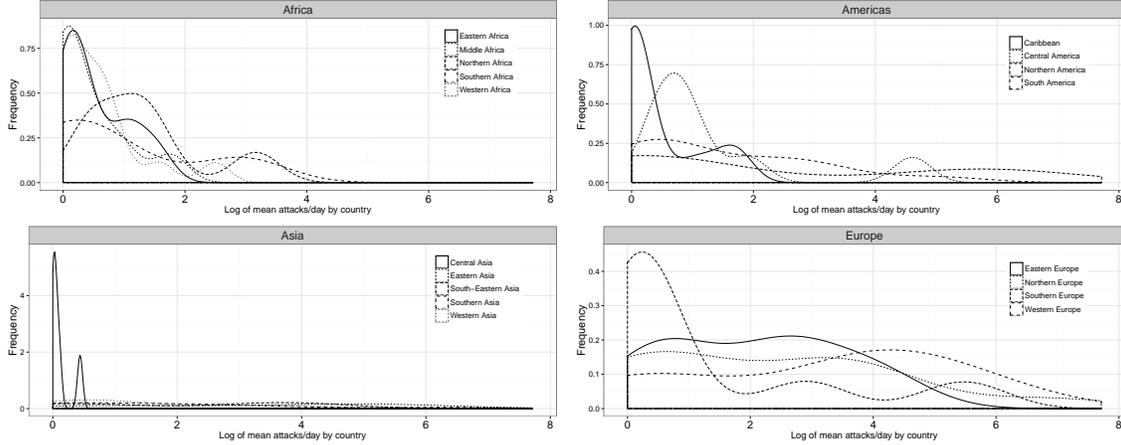


Figure 3: Distribution of attacks per day received in different geographic regions.

Note: Kernel density of mean cumulative attacks per day by geographical region. Regions in Americas, Asia, and Europe show the highest rates of attacks. Attack densities vary can vary substantially per region (possibly also depending on the vendor’s market share). Oceania accounts for a negligible fraction of attacks overall.

Table 4: Descriptive statistics of variables for \mathcal{T} , \mathcal{N} , \mathcal{U} , Hst, Frq, Pk, Geo.

Variable	Mean/Freq.	St. dev.	Obs.	Variable	Freq.	Obs.
Delay, Volume, Machines				Geo		
\mathcal{T}	0.99	0.831	2.57	Australia & New Zeal.	0.039	0.101
\mathcal{N}	13.552	102.28	2.57	Caribbean	0.013	0.033
\mathcal{U}	11.965	93.11	2.57	Central America	0.007	0.018
Hst				Central Asia	0	0.000063
Ev	0.552		1.42	Eastern Africa	0.001	0.00173
Rm	0.109		0.280	Eastern Asia	0.043	0.109
St	0.076		0.195	Eastern Europe	0.011	0.0282
Up	0.263		0.677	Melanesia	0	0.000159
Frq				Micronesia	0.001	0.00312
XH	0.004		0.0101	Middle Africa	0	0.000112
H	0.179		0.461	Not Avail.	0.112	0.289
L	0.436		1.12	Northern Africa	0.002	0.00409
M	0.379		0.975	Northern America	0.468	1.20
VH	0.001		0.0357	Northern Europe	0.023	0.0600
Pk				Polynesia	0	0.0000028
XH	0		0.000292	South America	0.008	0.0208
H	0.296		0.760	South-Eastern Asia	0.023	0.0593
L	0.091		0.235	Southern Africa	0	0.00103
M	0.609		1.57	Southern Asia	0.013	0.0339
VH	0.004		0.01018	Southern Europe	0.063	0.161
				Western Africa	0.001	0.00296
				Western Asia	0.018	0.0468
				Western Europe	0.154	0.395

To evaluate these scenarios we identify three types of attack pairs that are summarized in Table 6: in the first type of attack pair (A_1) the first attacks and the second attack affect the same vulnerability and, consequently, the same software version; in the second pair (A_2) the first attack and the second attack affect the same software, but different CVEs and different software versions; finally the first and second attacks affect the same software and the same version but

Table 5: Descriptive statistics for CVE₁ and CVE₂ variables.

CVE ₁			CVE ₂		
Variable	Mean	Obs.	Variable	Mean	Obs.
Compl _{CVE1,H}	0.009	23K	Compl _{CVE2,H}	0.009	24K
Compl _{CVE1,L}	0.42	1.08M	Compl _{CVE2,L}	0.334	0.86M
Compl _{CVE1,M}	0.571	1.47M	Compl _{CVE2,M}	0.657	1.69M
Imp _{CVE1}	9.549	2.57M	Imp _{CVE2}	9.681	2.57M
Internet Explorer	0.096	.25M	Internet Explorer	0.04	0.10M
PLUGIN	0.791	2.03M	PLUGIN	0.9	2.31M
PROD	0.083	.21M	PROD	0.037	95K
SERVER	0.03	77K	SERVER	0.023	59K
Pub. Year	2008.8	2.57M	Pub. Year	2009.4	2.57M

Table 6: Sample Attack Scenarios and Compatibility with Work-Aversion Hypothesis

Type	Condition	Description	Hypothesis
A ₁	CVE ₁ = CVE ₂	The first attacks and the second attack affect precisely the same vulnerability and, consequently, the same software version	Often for Hyp 3 as $T^* \rightarrow \infty$
A ₂	CVE ₁ ≠ CVE ₂ ∧ Sw _{CVE₁} = Sw _{CVE₂} ∧ Ver _{CVE₁} ≠ Ver _{CVE₂}	The first attack and the second attack affect the same software but different CVEs and different software versions.	Less frequent for Hyp 1 and Hyp 2 as $0 < T^* < \infty$
A ₃	CVE ₁ ≠ CVE ₂ ∧ Sw _{CVE₁} = Sw _{CVE₂} ∧ Ver _{CVE₁} = Ver _{CVE₂}	First and second attacks affect the same software and the same version but exploit different vulnerabilities	Almost never for Hyp 1 as $\theta_{V \cup \{v\}} = \theta_V$

Note: We expect the vast majority of attacks generated by the work-averse attacker to be of type A₁. A₂ should be less frequent than A₁, as it requires to engineer a new exploit. A₃ contradicts the work aversion hypothesis and should be the least common type of attack.

exploit different vulnerabilities (A₃). According to our hypothesis we expect that A₁ should be more popular than A₂ (in particular when the delay between the attacks is small) whilst A₃ should be the least popular of the three. To evaluate these attacks it is important to consider that users have diverging models of software security (Wash, 2010), different software have different update patterns, frequencies and attack vectors, see Nappa et al. (2015), and Provos et al. (2008).

For example, an attack against a browser may only require the user to visit a webpage, while an attack against a word processing application may need the user to actively open a file on the system (see also the definition of the **Attack Vector** metric in the CVSS standard CVSS-SIG (2015)). As these clearly require a different attack process, we further classify Sw in four categories: SERVER, PLUGIN, PROD(-ductivity) and Internet Explorer. The categories are defined by the software names in the database. For example SERVER environments are typically better maintained than ‘consumer’ environments and are often protected by perimetric defenses such as firewalls or IDSs. This may in turn affect an attacker’s attitude toward developing new exploits. This may require the attacker to engineer different attacks for the same software version in order to evade the additional mitigating controls in place. Hence we expect the difference between A₂ and A₃ to be narrower for the SERVER category.

Figure 4 reports a fitted curve of targeted machines as a function of time by software category. As expected, A₁ dominates in all software types. The predicted order is valid for PLUGIN and PROD. For PROD software we find no attacks against new vulnerabilities for different software

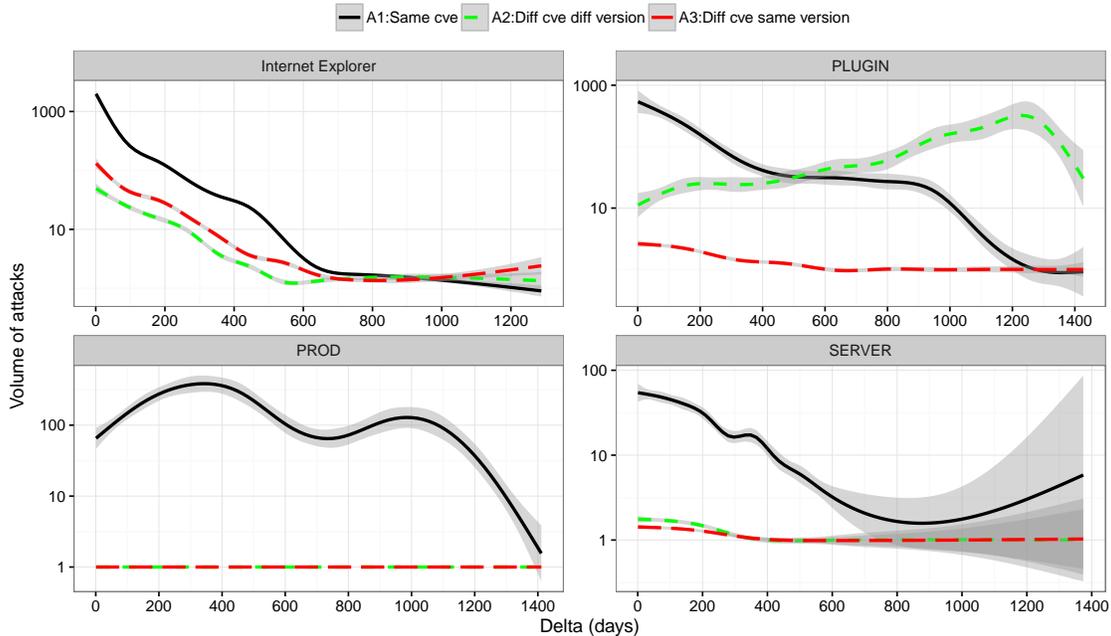


Figure 4: Loess regression of volume of attacks in time.

Volume of received attacks as a function of time for the three types of attack. A_1 is represented by a solid black line; A_2 by a long-dashed red line; A_3 by a dashed green line. The grey areas represent 95% confidence intervals. For Internet Explorer vulnerabilities the maximum \mathcal{T} between two attacks is 1288 days; for SERVER is 1374 days; PROD 1411; PLUGIN 1428. This can be determined by the timing of first appearance of the attack in the WINE database.

versions, therefore $A_2 = A_3 = 0$. This may be an effect of the low update rate of this type of software and relatively short timeframe considered in our dataset (3 years), or of a scarce attacker interest in this software type. Results for SERVER are mixed: the difference between A_2 and A_3 is very narrow and A_3 is occasionally higher than A_2 . Since oscillations occur within confidence intervals they might be due to chance.

Internet Explorer is an interesting case in itself. Here, contrary to our prediction, A_3 is higher than A_2 . By further investigating the data, we find that the reversed trend is explained by one single outlier pair: $CVE_1 = CVE-2010-0806$ and $CVE_2 = CVE-2009-3672$. These vulnerabilities affect Internet Explorer version 7 and have been disclosed 98 days apart, within our 120 days threshold. More interestingly, they are very similar: they both affect a memory corruption bug in Internet Explorer 7 that allows for a heap-spray attack resulting in arbitrary code execution. Two observations are particularly interesting:

1. Heap spray attacks are unreliable attacks that may result in a significant drop in exploitation success. This is reflected in the `Access Complexity:Medium` assessment assigned to both vulnerabilities by the CVSS v2 framework. In our model, this would imply a lower return $r(t, N_V(t), V)$ for the attacker, as the unreliable exploit may yield control of fewer machines among the vulnerable ones.

2. The exploitation code found on Exploit-DB⁹ is essentially the same for these two vulnerabilities. The code for CVE₂ is effectively a rearrangement of the code for CVE₁, with different variable names. In our model, this would indicate that the cost $C(v|V) \approx 0$ to build an exploit for the second vulnerability is negligible, as most of the exploitation code can be re-used from CVE₁ (See Appendix §A for details).

Hence, this vulnerability pair is only an apparent exception: the very nature of the second exploit for Internet Explorer 7 is coherent with our model and in line with Hyp. 1 and Hyp. 2. Removing the pair from the data confirms the order of attack scenarios identified in Table 6.

We now check how the trends of attacks against a software change with time. Hyp. 3 states that the exploitation of the same vulnerability persists in time and decreases slowly at a pace depending on users’ update behaviour. This hypothesis offers an alternative behavior with respect to other models in literature where new exploits arrive very quickly after the date of disclosure, and attacks increase following a steep curve as discussed by Arora et al. (2004).

5.1 An Econometric Model of the Engineering of Exploits

We can use Proposition 2 to identify a number of additional hypothesis that are useful to formulate the regression equation. At first we notice that $T^* = O(\log(\theta_v - \theta_V)N)$. Therefore we have a first identification relation between the empirical variable \mathcal{U} (corresponding to N) and the empirical variable \mathcal{T} (whose correspondence to T^* is outlined later in this section).

Hypothesis 4. *The relation between number of attacked systems \mathcal{U} and delay \mathcal{T} is log-linear.*

Since $\partial T^*/\partial((\theta_v - \theta_V)N) < 0$ a larger number of attacked systems \mathcal{U} on *different* versions ($\theta_v \neq \theta_V$) would imply a lower delay \mathcal{T} (as there is an attractive number of new systems that guarantee the profitability of new attacks). In contrast, the baseline rate of attacks impacts negatively the optimal time \mathcal{T} as $\partial T^*/\partial(\theta_V N) > 0$ since a larger pool of vulnerable machines makes it more profitable to continue with existing attacks (as per Hyp. 1).

Hypothesis 5. *The possibility of launching a large number of attacks against systems for which an exploit already exists lengthens the time for weaponizing a vulnerability ($\mathcal{N} \cdot (\text{Ver}_0 = \text{Ver}_v) \uparrow \implies \mathcal{T} \uparrow$), whereas an increase in potential attacks on different systems is an incentive towards a shorter weaponization cycle ($\mathcal{N} \cdot (\text{Ver}_0 \neq \text{Ver}_v) \uparrow \implies \mathcal{T} \downarrow$).*

When considering the effects of costs, we observe that, as $\partial T^*/\partial C(v|V) > 0$, the presence of a vulnerability with a low attack complexity implies $dC(v|V) < 0$, and therefore reflects a drop in the delay \mathcal{T} between the two attacks. We have already discussed this possibility as Hypothesis 2.

As for revenues, it is $\partial T^*/\partial r < 0$ so that an higher expected profit would imply a shorter time to weaponization. However, we cannot exactly capture the latter condition since in our model the actual revenue r is stationary and depends only on the captured machine rather than the individual type of exploit. A possible proxy is available through the impact variable Imp , but it only shows the level of *technical* compromise that is possible to achieve. Unfortunately, such information might not correspond to the actual revenue that can be extracted by the attacker. For example, vulnerabilities that only compromise the availability of a system are scored low according to the CVSS standard. However, for an hacker offering “booter services” to on-line gamers (i.e. DDoS targeted attack

⁹See Exploit-DB (<http://www.exploit-db.com>, last accessed January 15, 2017.), which is a public dataset for vulnerability proof-of-concept exploits. CVE₁ corresponds to exploit 16547 and CVE₂ corresponds to exploit 11683.

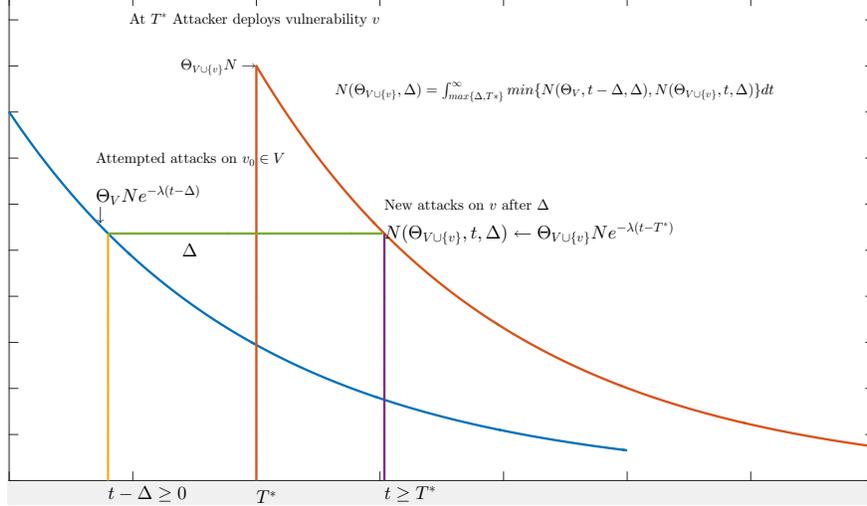


Figure 5: Computing the delay (\mathcal{T}) against different vulnerabilities.

Change in the number of attacked systems for two attacks against different systems $\Delta = \mathcal{T}$ days apart. The first attack happens at $t - \mathcal{T} \geq 0$ and the number of attacked systems $\mathcal{U}(\Theta_V \cup \{v\}, t, \mathcal{T})$ is derived from Eq. (1) as $\Theta_V N e^{-\lambda(t-\mathcal{T})}$. The number of systems attacked by the new exploit introduced at T^* is derived as $\mathcal{U}(\Theta_{V \cup \{v\}}, t, T^*) = N \Theta_{V \cup \{v\}} e^{-\lambda(t-T^*)} dt$.

against fellow gamers) these vulnerabilities are the only interesting source of revenues Hutchings and Clayton (2016).

However **Imp** can also be seen as a potential conditional factor to boost the attractiveness of a vulnerability as the additional costs of introducing an exploit might be justified by the increased capability to produce more problems for the target if this is the objective of the attacker.

Hypothesis 6. *Vulnerabilities with higher impact increase revenue and therefore decrease number of attacks ($\text{Imp}_{\text{CVE}_2} > \text{Imp}_{\text{CVE}_1} \implies \mathcal{U} \downarrow$).*

As the time of introduction of an exploit T^* can not be directly measured from our dataset, we use \mathcal{T} (i.e. the time in between two consequent attacks) as a proxy for the same variable. Figure 5 reports a pictorial representation of the transformation. Each curve represents the decay in time of number of attacks against two different vulnerabilities. The first attack (blue line) is introduced at $t = 0$, and the second (red line) at $t = T^*$. The number of received attacks is described by the area below the curve within a certain interval. Let $\mathcal{U}(\Theta_V \cup \{v\}, t, \mathcal{T})$ represent the number of systems that receive two attacks \mathcal{T} days apart, at times $t - \mathcal{T}$ and t respectively. Depending on the relative position of $t - \mathcal{T}$ with respect to T^* , the interval within which the attacks on the pair of vulnerabilities will be measured is $\int_{\max(T^*, \mathcal{T})}^{\infty} \mathcal{U}(\cdot) dt$. Setting the number of attacks at time $t - \mathcal{T}$ as $\mathcal{U}(\theta_v, t - \mathcal{T}) = N \theta_v e^{-\lambda(t-\mathcal{T})}$ and the attacks received on the second vulnerability at time t as $\mathcal{U}(\theta_{V \cup \{v\}}, t) = N \theta_{V \cup \{v\}} e^{-\lambda(t-T^*)}$, we obtain

$$\mathcal{U}(\theta_{V \cup \{v\}}, t, \mathcal{T}) = \min \left(N \theta_v e^{\lambda \mathcal{T}}, N \theta_{V \cup \{v\}} e^{\lambda T^*} \right) \int_{\max(\mathcal{T}, T^*)}^{\infty} e^{-\lambda t} dt \quad (9)$$

Solving for the two cases $T^* > \mathcal{T}$ and $T^* < \mathcal{T}$, we formulate the following claim:

Claim 1. *The sign of the coefficient for \mathcal{T} oscillates from positive to negative as \mathcal{T} increases.*

$$\log \mathcal{U}(\theta_{V \cup \{v\}}, t, \mathcal{T}) = \begin{cases} \log \frac{N}{\lambda} - \lambda T^* + \lambda \mathcal{T} + \log \theta_v & \text{if } T^* > \mathcal{T} \\ \log \frac{N}{\lambda} + \lambda T^* - \lambda \mathcal{T} + \log \theta_{V \cup \{v\}} & \text{if } T^* < \mathcal{T} \end{cases} \quad (10)$$

Proof of Claim 1 and its empirical evaluation are given in Appendix 7.4. \square

Figure 1 indicates that \mathcal{T} is on average more than 100 days with respect to T^* , as such:

$$\log \mathcal{U} = -\lambda \mathcal{T} + \lambda T^* + \log \frac{N}{\lambda} + \log \theta_{V \cup \{v\}}$$

Substituting T^* from Eq. (8), the number of expected attacked systems after \mathcal{T} days is:

$$\log \mathcal{U} = -\lambda \mathcal{T} + \lambda \left[\frac{1}{\delta} \log \left(\frac{C(v|V)}{r} - \frac{\delta}{\lambda + \delta} (\theta_{V \cup \{v\}} - \theta_V) N \right) \right] + \log \frac{N}{\lambda} + \log \theta_{V \cup \{v\}}. \quad (11)$$

Our regression model tests the hypotheses above by reflecting the formulation provided in Eq. (11). \mathcal{T} can be measured directly in our dataset; the cost of development of an exploits $C(v|V)$ can be estimated by the proxy variables $\text{Comp}_{\text{CVE}_2}$, as the complexity associated with exploit development requires additional engineering effort (and is thus related to an increase in development effort) CVSS-SIG (2015).

We can not directly measure the revenue r and the number of systems N affected by the vulnerability, but we can estimate the effect of an attack on a population of users by measuring the impact (Imp) of that vulnerability on the system: higher impact vulnerabilities (i.e. ($\text{Imp}_{\text{CVE}_2} > \text{Imp}_{\text{CVE}_1}$) allow the attacker to control a *higher fraction* of the vulnerable system, and therefore extract higher revenue r from the attack. Similarly, the introduction of an attack with a higher impact can approximate the difference in attack penetration $(\theta_{V \cup \{v\}} - \theta_V)N$ for the new set of exploits as it allows the attacker for a higher degree of control on the affected systems. Finally, high impact vulnerabilities ($\text{Imp}_{\text{CVE}_2, H}$), for example allowing remote execution of arbitrary code on the victim system, leave the $\Theta_{V \cup \{v\}}N$ systems under complete control of the attacker; in contrast, a low impact vulnerability, for example causing a denial of service, would allow for only a temporary effect on the machine and therefore a lower degree of control.

In Table 7 we report the sample correlation matrix for the variables included in the regression system we will use to parameterize the model, from an econometric standpoint, the highest pairwise correlations with \mathcal{T}_i are **Frq** MEDIUM and **Pk** HIGH, however these have correlations of less than 20%, as such the standard issues on rank and collinearity are not present.

The model predicts that for each pair the logarithm of \mathcal{U}_i should be statistically related to \mathcal{T}_i with the coefficient parameterizing attackers anticipated arrival of patches and antivirus signatures λ , presuming the additional controls orthogonalize the effort effects from the complexity and opportunity set provided by the attack type. The residual of the regression can then be interpreted as being the function T^* of the idiosyncratic stochastic discount factor δ that each individual attacker applies when making their decision to refresh their malware technology.

As we cannot ascribe a specific collection of observations to a specific individual, this is an unobserved variable. To attempt to correct for this statistical feature, we attempt to identify commonalities in attacks by including the significant principal components estimated from user data and geographical location, see Figure 6 and related discussion for details.

As such, we estimate three equations to evaluate the effect of our regressors on the dependent

Table 7: Correlation Matrix of All Variables Included in the Model.

Model variable	1.	2.	3.	4.	5.	6.	7.	8.	9.	10.	11.
1. \mathcal{T}	1										
2. Comp1_{CVE2L}	-0.130	1									
3. Imp_{CVE2H}	0.058	-0.237	1								
4. $\text{Imp}_{CVE2} > \text{Imp}_{CVE1}$	0.092	-0.097	0.055	1							
5. Geo North. Am.	0.020	0.146	-0.040	-0.035	1						
6. Geo Western Eu.	0.022	-0.026	0.011	-0.051	-0.410	1					
7. Hst EVOLVE	-0.084	0.024	-0.065	0.022	-0.057	0.015	1				
8. Hst UPGRADE	0.054	0.004	0.030	-0.017	0.084	-0.042	-0.656	1			
9. Frq HIGH	0.104	-0.050	0.052	0.107	0.087	-0.183	-0.225	-0.024	1		
10. Frq MEDIUM	0.136	0.077	-0.103	0.127	0.014	0.096	-0.087	0.186	-0.369	1	
11. Pk HIGH	0.166	-0.008	0.045	0.049	0.020	0.033	-0.288	0.034	0.678	0.052	1
12. Pk MEDIUM	-0.087	0.031	-0.054	-0.004	-0.023	-0.008	0.197	0.004	-0.549	0.101	-0.814

Notes: The correlation matrix is presented for the whole sample, for variable definitions see Table 1.

Table 8: Summary of predictions derived from the model.

Model variable	Regressor	Expectation	Hyp.	Rationale
\mathcal{T}	\mathcal{T}	$\beta_1 < 0$	Hyp. 3, Hyp. 4, Hyp. 5	Shorter exploitation times are associated with more vulnerable systems, hence $\mathcal{T} \uparrow \implies \mathcal{U} \downarrow$.
$C(V v)$	$\text{Comp1}_{CVE2,L}$	$\beta_2 < 0$	Hyp. 1, Hyp. 5, Hyp. 2	The introduction of a new reliable, low-complexity exploit minimizes implementation costs, thus $C \downarrow \implies \mathcal{U} \downarrow$.
$\theta_{V \cup \{v\}}$	$\text{Imp}_{CVE2,H}$	$\beta_3 > 0$	Hyp. 6, Hyp. 5	High impact vulnerabilities allow the attacker for a complete control of the attacked systems, hence $\theta_{V \cup \{v\}} \uparrow \implies \mathcal{U} \uparrow$.
$r, (\theta_{V \cup \{v\}} - \theta_V)$	$\text{Imp}_{CVE2} > \text{Imp}_{CVE1}$	$\beta_4 < 0$	Hyp. 6	Selecting a higher impact exploit for a new vulnerability increases the expected revenue and increases the fraction of newly controlled systems with respect to the old vulnerability. $r \uparrow \implies \mathcal{U} \downarrow$ and $(\theta_{V \cup \{v\}} - \theta_V) \uparrow \implies \mathcal{U} \downarrow$.

variable. The formulation is derived from first principles in Eq. (11). Our equations are:

$$\text{Model 1: } \log(\mathcal{U}_i) = \beta_0 + \beta_1 \mathcal{T}_i + \epsilon_i \quad (12)$$

$$\text{Model 2: } \log(\mathcal{U}_i) = \dots + \beta_2 \text{Comp1}_{i,CVE2,L} + \epsilon_i \quad (13)$$

$$\text{Model 3: } \log(\mathcal{U}_i) = \dots + \beta_3 \text{Imp}_{i,CVE2,H} + \beta_4 (\text{Imp}_{i,CVE2} > \text{Imp}_{i,CVE1}) \epsilon_i \quad (14)$$

Where i indexes the pair of attacks received by each machine after \mathcal{T} days, $\text{Comp1}_{i,CVE2,L}$ indicates that CVE_2 has a low complexity, and $\text{Imp}_{i,CVE2,H}$ indicates that CVE_2 has a *High* (≥ 7) impact. Both classifications for Comp1 and Imp are reported by the CVSS standard specification. The correlation matrix of the model variables is reported in Table 7.

The mapping of each term with our hypotheses and the predicted values of the regressors are described in the Table 8. Further, we add the vector of controls Z to the regression to account for exogenous factors that may confound the observation, as discussed in Section 4.2: Geo , Hst , Frq , Pk . Z is defined as

$$Z_i^{\text{all}} = \sum_{d \in \text{Geo}} \alpha_d (\text{Geo}_i = d) + \sum_{d \in \text{Hst}} \alpha_d (\text{Hst}_i = d) + \sum_{d \in \text{Frq}} \alpha_d (\text{Frq}_i = d) + \sum_{d \in \text{Pk}} \alpha_d (\text{Pk}_i = d) \quad (15)$$

Given the model, it is reasonable to expect the variance in attacked systems to be related to the level of some independent variables. For example, the variance of the dependent variable \mathcal{U} may depend

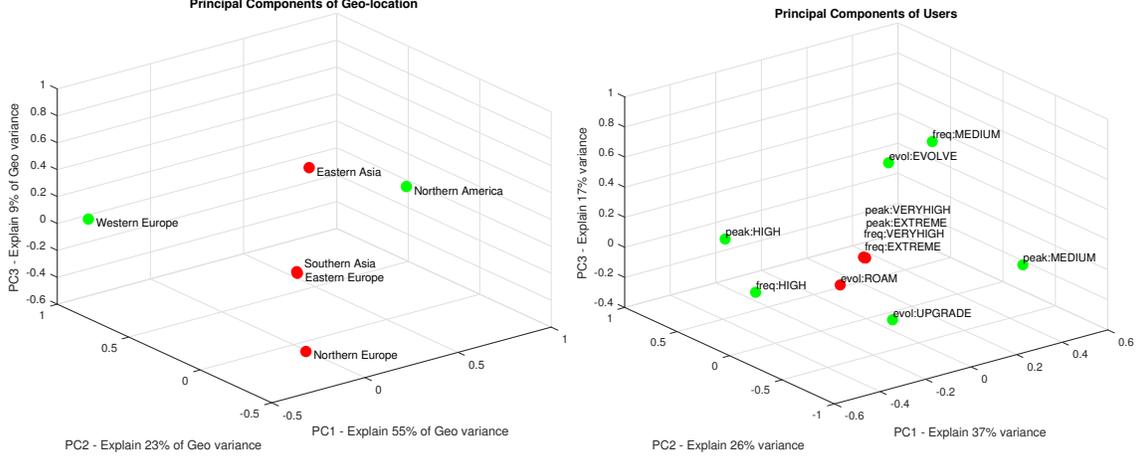


Figure 6: Principal component analysis of location and user factors.

Each dot corresponds to the loading values for the geographical location and user specific factors denoted **Geo**, **Hst**, **Frq**, and **Pk** factors. The axes report the loadings of the components for the eigenvalues that explain the most variance. The farthest points from the graph origin (in green) are factor levels with the highest level of independence.

on the value of \mathcal{T} . To address possible heteroskedasticity of the data, we employ a robust regression next to a regular OLS and compare the results. To evaluate the independence between the control factor levels over the respective contingency table we employ a standard Principal Component Analysis. Figure 6 shows a plot of the analysis for loadings of the principal components in Table 2.

We select all eigenvalues that explain at least 10% of the variance: **Geo** has only three eigenvalues above the threshold, whereas **Hst**, **Frq**, and **Pk** have four eigenvalues above the threshold. Altogether they explain 87% of the variance for geo-location and 80% for user characteristics. The corresponding components of the eigenvectors (corresponding to the loadings of the controls) are identified and we select the control values that have the greatest distance from the origin in the eigenvector space (at least 10% again as a sum of squares). This guarantees that we select the controls with the highest degree of independence. This results in eight selected controls, identified in green in Figure 6.

Our regression results are reported in Table 9. We utilize two estimators as we have little information on the error structure of the regression model and we are subject to certain statistical issues caused by the right truncation of the data, that is we do not observe \mathcal{T} asymptotically by construction.

First is a simple OLS estimator with Huber-White standard errors and second is a Robust fit model that utilizes a WLS type estimator with iterative re-weighting and we implement the sandwich form standard error from the WLS iterations. The weighting function for the iterative re-weighting is a bisquare function, experimentation with spectral and Andrews type weightings suggest the regressions are insensitive to kernel and tuning function. For the robust fit we compute a McFadden adjusted pseudo- R^2 , which sets the numerator as the log likelihood function at the estimate and the denominator as the log likelihood of just the intercept alone. Note that it is not appropriate to compare directly the pseudo- R^2 and the R^2 from the OLS estimates, which suggests that the model captures roughly 10% of the variation in numbers of attacked machines, as opposed to explaining 35% of the model likelihood for the pseudo- R^2 .

The set of OLS and Robust regressions returns very similar estimations. We also experimented with various regression estimators (e.g. 2SLS, 3SLS) and they produced markedly similar results to OLS, subject to the standard caveats on mis-identification. The introduction of the controls only change the sign of β_1 from positive to negative for Model 1. This may indicate that the type of user is a significant factor in determining the number of delivered attacks, which is consistent with previous findings Nappa et al. (2015). Interestingly, the factor that introduces the highest change in the estimated coefficient β_1 for \mathcal{T} is `Comp1` (Model 2), whereas its estimate remains essentially unchanged in Model 3. This may indicate that the cost of introduction of an exploit has a direct impact on the time of delivery of the exploit. The coefficients for all other regressors are consistent across models, and their magnitude changes only slightly with the introduction of the controls. This observation is to be expected: user characteristics should not influence the characteristics of the vulnerabilities present on the system; as such, the distribution of attacks in the wild seems to depend mostly on system characteristics rather than user type.

The signs of coefficients for the `Imp` variables suggest that both impact of a new vulnerability and its relation with the impact of previous vulnerabilities have an effect on the number of attacked systems. Interestingly, a *high* impact encourages the deployment of attacks and increases the number of attacked systems, whereas the introduction of a *higher* impact vulnerability requires the infection of a smaller number of systems as revenues extracted from each machine increase. Hence, when introducing a new exploit, the attacker will preferably choose one that grants a higher control over the population of users ($\theta_{V \cup \{v\}} > \theta_V$) and use it against a large number of system. This is consistent with recent findings suggesting that vulnerability severity alone is not a good predictor for exploitation in the wild Allodi and Massacci (2014); Bozorgi et al. (2010). Other factors such as software popularity may play a role Nayak et al. (2014).

6 Discussion, Conclusions and Implications

This paper implements a model of the *Work-Averse Attacker* as a new conceptual framing to understand cyber threats. Our model presumes that an attacker is a resource-limited actor with fixed costs that has to choose which vulnerabilities to exploit to attack the ‘mass of Internet systems’. Work aversion simply means that effort for the attacker is costly (in terms of cognition and opportunity costs), hence a trade-off exists between effort exerted on new attacking technologies and the anticipated reward schedule from these technologies. As technology combinations mature, their revenue streams are presumed dwindle.

In this framework, an utility-maximizing attacker will drive exploit production according to their expectations that the newly engineered attack will increase net profits from attacks against the general population of internet users. As systems in the wild get patched unevenly and often slowly in time (Nappa et al., 2015), we model the production of new vulnerability exploits following Stokey’s ‘economy of inaction’ logic, whereby ‘doing nothing’ before a certain (time) threshold is the best strategy. A cost constraint driving the attacker’s exploit selection strategy naturally emerges from the model. In particular, we find theoretical and empirical evidence that:

1. An attacker massively deploys only one exploit per software version. The only exception we found is for Internet Explorer; the exception is characterized by a very low cost to create an additional exploit, where it is sufficient to essentially copy and paste code from the old exploit, with only few modifications, to obtain the new one. This finding is predicted by the model and supports Hyp. 1.

Table 9: Ordinary Least Squares and Robust Regression Results

Dependent Variable: natural logarithm of the number of attacked machines $\log(\mathcal{U}_i)$												
	Model 1				Model 2				Model 3			
	OLS Z1 : Z8		Robust Z1 : Z8		OLS Z1 : Z8		Robust Z1 : Z8		OLS Z1 : Z8		Robust Z1 : Z8	
c	0.927	0.006	0.731	0.096	1.065	0.122	0.845	0.171	0.933	-0.106	0.783	0.039
\mathcal{T}	(0.001)	(0.003)	(0.001)	(0.003)	(0.001)	(0.003)	(0.001)	(0.003)	(0.004)	(0.005)	(0.003)	(0.004)
Compl _{CVE2 L}	0.018	-0.051	0.012	-0.044	-0.006	-0.092	-0.003	-0.071	-0.005	-0.091	-0.004	-0.071
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Imp _{CVE2 H}									0.144	0.236	0.063	0.131
									(0.003)	(0.003)	(0.003)	(0.003)
Imp _{CVE2 > Imp_{CVE1}}									-0.088	-0.209	0.012	-0.087
									(0.003)	(0.003)	(0.002)	(0.002)
Z1: Geo North. Amer.	0.604		0.37		0.679		0.422		0.671		0.419	
	(0.002)		(0.001)		(0.002)		(0.001)		(0.002)		(0.001)	
Z2: Geo West. Eu.	0.155		0.105		0.17		0.116		0.163		0.114	
	(0.002)		(0.002)		(0.002)		(0.002)		(0.002)		(0.002)	
Z3: Hst EVOLVE	0.191		0.129		0.208		0.141		0.223		0.149	
	(0.002)		(0.002)		(0.002)		(0.002)		(0.002)		(0.002)	
Z4: Hst UPGRADE	0.112		0.072		0.116		0.076		0.113		0.075	
	(0.002)		(0.002)		(0.002)		(0.002)		(0.002)		(0.002)	
Z5: Frq HIGH	0.24		0.147		0.212		0.127		0.279		0.157	
	(0.003)		(0.003)		(0.003)		(0.003)		(0.003)		(0.003)	
Z6: Frq MEDIUM	0.328		0.227		0.358		0.246		0.41		0.271	
	(0.002)		(0.002)		(0.002)		(0.002)		(0.002)		(0.002)	
Z7: Pk HIGH	0.513		0.442		0.567		0.49		0.531		0.477	
	(0.004)		(0.003)		(0.004)		(0.003)		(0.004)		(0.003)	
Z8: Pk MEDIUM	0.379		0.274		0.412		0.299		0.411		0.301	
	(0.003)		(0.002)		(0.003)		(0.002)		(0.003)		(0.002)	
Pseudo R^2	-	-	0.326	0.341	-	-	0.331	0.347	-	-	0.331	0.347
R^2	0.00	0.093	-	-	0.016	0.126	-	-	0.017	0.13	-	-
F	348.66	26,551.47	-	-	18,548.25	33,422.78	-	-	9,989.88	28,915.60	-	-
Obs.	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500	2,324,500

$$\text{Model 1: } \log(\mathcal{U}_i) = \beta_0 + \beta_1 \mathcal{T}_i + \epsilon_i$$

$$\text{Model 2: } \log(\mathcal{U}_i) = \beta_0 + \beta_1 \mathcal{T}_i + \beta_2 \text{Compl}_{i,CVE2,L} + \epsilon_i$$

$$\text{Model 3: } \log(\mathcal{U}_i) = \beta_0 + \beta_1 \mathcal{T}_i + \beta_2 \text{Compl}_{i,CVE2,L} + \beta_3 \text{Imp}_{i,CVE2,H} + \beta_4 \text{Imp}_{i,CVE2 > Imp_{i,CVE1}} + \epsilon_i$$

Notes: The three model equations reflect the definition of the expected (log) number of affected machines after an interval \mathcal{T} . The regression model formulation is derived from prime principle from Eq. 11. The expected coefficient signs are given in Table 8. For each model we run four sets of regressions. OLS and Robust regressions are provided to address heteroscedasticity in the data. R^2 and F - statistics are reported for the OLS estimations. Note that the pseudo- R^2 are computed for the robust regressions, using the McFadden adjusted approach $R^2 = 1 - (\log(LL_{full}) - K) / \log(LL_{int})$, where $\log(LL_{full})$ is the log likelihood for the full model minus the number of slope parameters K versus the log likelihood of the intercept alone and should not be compared directly to the OLS R^2 . Coefficient estimations of the two sets of regressions are consistent. All regressions are run with and without the set of controls defined by the PCA analysis pictorially reported in Fig. 6.

All coefficient signs for the three models reflect the work-averse attacker model predictions, with the only exception of the estimation for \mathcal{T} with no controls for which the prediction for β_1 is inverted. This may indicate that user characteristics are relevant factors for the arrival time of exploits when other factors related to the system are not accounted for. The introduction of **Compl** in Model 2 significantly changes the estimate for β_1 , whereas **Imp** in Model 3 leaves the estimates for **Compl** and \mathcal{T} unchanged. High **Imp** vulnerabilities tend to increase volume of attacks. We report only standard errors without starring p-values as all coefficients are significant due to the number of observations in the dataset. All standard errors are estimated using the Huber-White approach.

2. Low complexity vulnerabilities for which a reliable exploit can be easily engineered lower the production costs and favor the deployment of the exploit. This finding supports Hyp. 2.
3. The attacker deploys new exploits relatively slowly over time, driven by a slowly decreasing instantaneous profit function; empirically, we find that attacks 1000 days apart are still driven by the same exploits in about 20% of the cases, and that the effect of the passage of time in between attacks (\mathcal{T}) on the number of affected system is indeed negative and very small given the current patching rate. This finding supports Hyp. 3 and Hyp. 5.
4. The presence of a high impact vulnerability increases the incidence of exploitation in the wild. Similarly, gaining a higher control over attacked systems heightens the attacker’s revenue and decreases the number of systems that need to be infected to balance costs. This supports Hyp. 6.

The above findings suggest that risk associated with software vulnerabilities can not be solely measured in terms of technical severity, and that other, measurable factors may be included in the assessment, as previously suggested by some authors (see for example Bozorgi et al., 2010; Houmb et al., 2010; Allodi and Massacci, 2014). While characteristics of attackers and users, such as skill or competence in system security, may remain hard to measure, this paper shows how measurable environmental indexes may be used to estimate economic incentives that exists for the attacker (as suggested, e.g. by Anderson, 2008).

6.1 Reconstructing the Attacker Discount Factor

Given the estimates in Table 9 the final step is to reconstruct the curves that relate the expected revenue per machine versus the minimum implied discount factor that an attacker would require to engage in those attacks. There are some caveats to this approach, first, the distribution of stochastic discount factors across attackers is unimodal and the distribution of machines attacked is stationary. Hence, solving Equation 11 by plugging in values of $\log \mathcal{U}$, $\lambda \mathcal{T}$, $\log \theta_{V \cup \{v\}}$ from the distributions observed in the dataset we can reconstruct the curves relating the discount rate δ to a grid of revenues r . As would be expected the discount rate is a downward sloping curve with respect to an increasing return on investment. However, we can see that for many feasible values of the expected return on investment per machine, the discount factor attains reasonably plausible values.

6.2 Limitations

Our model describes an attacking adversary with costly production of malware. The model aims at explaining the attacker’s preference in exploiting one vulnerability over another rather than casting their net repeatedly after each new vulnerability is discovered. This preference can not be directly measured, but must be inferred through their real attack signatures from a record of attempted security incursions.

Records of attacks detected over a user’s machine are necessarily conditioned over the user’s proneness in receiving a particular attack. For example, a user may be inclined to open executable email attachments, but not in visiting suspicious websites. Thus, there may be a disassociation between the observed attacks and those engineered by the attacker. For our empirical dataset this limitation is mitigated by *WINE* reporting attack data on a very large representative sample of Internet users (Dumitras and Shou, 2011). However, we also need to have additional conditioning

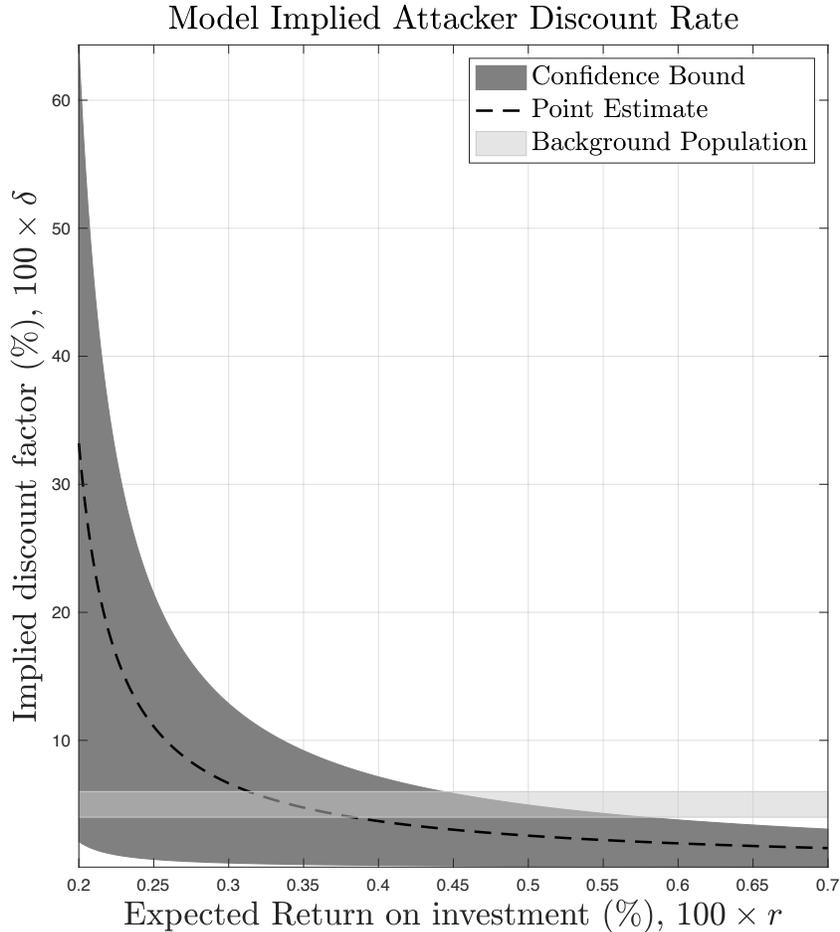


Figure 7: Discount rates versus expected return.

Implied relationship between the attacked discount rate δ and expected revenue per machine r .

variables to permit identification of the impact on various behavioral characteristics. Many of the additional characteristics of users that may influence the observed volume of attacks, such as educational level and culture which, are very difficult or close to impossible to gauge at the scale of data presented in this paper. As proxies to control for this effect we employ the **User Profile**, **Frequency**, **Peak** and geographic location variables, as these outline the user’s proneness in receiving attacks. Further, geographic location may not only influence effects related to user culture, but also on attack diffusion.

Software versioning information is known to be unreliable at times with respect to vulnerability existence (Nguyen et al., 2015). Further, software versions can not be easily ‘ordered’ throughout software types, as different vendors adopt different naming schemes for software releases (Christey and Martin, 2013, for an overview). We can not therefore order software versions over time easily. This is however irrelevant to our study as we are interested in measuring the sequences of newer

attacks received by internet users, as opposed to measuring the existence of new exploits for subsequent software releases, our model predicts that the attacker will perform this information filtration dynamically as they view the rewards from their activities. A limitation of our empirical dataset is obviously the market penetration of Symantec, as of 2016 Symantec self reports that it is the largest security vendor¹⁰ by market share in anti-virus and overall software security and hence has a broad coverage recording attacks on customers. However, third party verifiable measurement of these claims is difficult hence replication studies across different security vendors would be easily implemented, given the simplicity of our model specification

6.3 Implications For Theory

The modeling of cyber attacks and cyber risk has traditionally been centered on the vulnerabilities present on systems and their technical severity. Little theoretical discourse has been developed around the identification of models of vulnerability exploit production. Most assessments on whether the attacker will be successful are still produced, in practice, by means of ‘expert estimates’ (Wang et al., 2008) or ‘technical measures of vulnerabilities’ (Naaliel et al., 2014) that, however, are known *not* to correlate with attacker choices as shown by Bozorgi et al. (2010). That “some vulnerabilities are different than others” (Nayak et al., 2014) is now a known fact, yet a distinction between vulnerabilities is yet to be fully discussed. In recent years, this perspective has been reinforced as the limited nature of attacker resources became empirically apparent (Grier et al., 2012), and advances in security modeling started distinguishing between opportunistic and deliberate attacks (see Ransbotham and Mitra, 2009, for a discussion on this). Current studies modeling cybersecurity events typically think of the attacker as employing an undefined ‘attack generation function’, which often collapses to an ‘all-powerful’ attacker that may generate an attack for whichever vulnerability exists: as each attacker determines an effort function that directly identifies exploitable vulnerabilities not currently monetized by others, for a large enough number of attackers all vulnerabilities are covered and the maximum entropy condition is achieved. In this paper we demonstrate that there are common production obstacles in malware production: attackers will agglomerate around existing technologies and only upgrade when those technologies ceases to have greater marginal benefit than marginal cost.

This paper addresses fundamental assumptions on attacker strategies and parametrization by identifying and empirically validating the attack production function based on the notion of ‘work-averseness’ for the class of “*non-reactive*” (as opposed to “*reactive*”) attackers (van Dijk et al., 2013). This classification is implicitly present across the spectrum of research in the information security domain, for example referring to “*targeted*” or “*untargeted*” attackers (Laszka et al., 2013). Modeling decisions under these classes are subject to a number of assumptions on the strategies adopted by both attackers and defenders. For example, *period renewal* strategies consider a process of periodic evolution of attacker (and/or defender) capabilities under the “non-reactive” assumption; van Dijk et al. (2013) find that periodic strategies are strongly dominant under a number of variants of their game, even when the attacker can evaluate some parameter of the defender’s strategy (e.g. the defender’s renewal period). A number of questions remain however open: what is the process that characterizes the attacker’s renewal process? Can we identify contextual and/or technical considerations that affect periodicity? Are all attacks equivalent under the same strategy?

¹⁰Indeed, according to the 2016 market share Symantec has held this position for the last 15 years, see <https://www.symantec.com/en/uk/about/newsroom/analyst/reports>

Our work is of particular relevance here on two aspects: firstly, we identify a theoretical underpinning to the attacker’s periodic renewal process that extends current assumptions on its nature (e.g. technical severity vs exploit availability). This overcomes generalist assumptions on the factors characterizing this process by identifying aspects of malware engineering (e.g. vulnerability complexity or presence of ‘equivalent’ exploit) that affect the attacker’s renewal period, opening to further differentiations in cost-effective defender strategies (for example to the development of fine-grained models for efficient exploit patching selection on the defender side). Second, our theoretical and empirical findings strongly support previous results on the dominance of periodic, non-reactive strategies on the attacker side in settings where the defender is not reactive (van Dijk et al., 2013); this dominance remains even in the case where the attacker can reliably evaluate some characteristics of the defender’s setting (e.g. a system configuration, or an average patching rate as evaluated in Nappa et al. (2015)); this is a case of prominent importance in reality, as attacking tools and malware at large rely on an infrastructure that automatically optimizes attack delivery on the basis of target characteristics such as geographical location or system configuration (Grier et al., 2012; Kotov and Massacci, 2013). Similar settings are considered in the extant literature on attacker decision making and target selection (Mitra and Ransbotham, 2015), which our work extends by further characterizing the attack deployment process in terms of vulnerability and temporal properties.

This can in turn inform the debate on regulatory intervention to devise regulatory models that, on average, will increase the development costs for the attacker or incentivize the deployment of defensive measures. For example, recent developments presented in Zhao et al. (2017) go in the direction of evaluating policies for reward mechanisms in vulnerability finding (e.g. in so-called ‘bug bounty programs’); following our work, these may be extended to account for the periodicity of vulnerability exploitation, and optimize incentives for vulnerabilities of higher relevance for the attacker accordingly. Even targeted attacks require fixed costs, but it is still unclear whether such attacks could be captured by variation in the work aversion (e.g. by making reward a function of costs such as reconnaissance (Verizon, 2011; Bilge and Dumitras, 2012)) and whether these costs may affect the dominant strategies on the defender side (van Dijk et al., 2013).

6.4 Implications For Information Security Management and Policy

Our findings suggest that the rationale behind vulnerability exploitation could be leveraged by defenders to deploy more cost-effective security countermeasures. For example, it is well known that software updates correspond to an increased risk of service disruption (e.g. for incompatibility problems or updated/deprecated libraries). However, if most of the risk for a particular software version comes from a specific vulnerability, than countermeasures other than patching may be more cost-effective. For example, maintaining network IDS signatures might be a better option than updating the software, because one IDS signature eliminate the majority of risks faced by that system; a software patch may ‘overdo-it’ by fixing more vulnerabilities than necessary.

Further, this view on the attacker has repercussions on the impact of vulnerability disclosure in terms of security of the user or organization. Work in this direction explored both the economic incentives for sharing information security (Gal-Or and Ghose, 2005; Ransbotham et al., 2012) and the impact of vulnerability disclosure on attacks and firm value (Arora et al., 2008; Telang and Wattal, 2007). Some of this discussion resulted in recent open debates regarding policies for vulnerability disclosure, and the drafting of ISO standards to guide vulnerability communication to the vendor and to the public (e.g. ISO/IEC 29147:2014). For example, the United States

Department of Commerce NTIA forum for vendors, industry players, and security researchers to discuss procedures and timings of the vulnerability disclosure process.¹¹ This discussion is not currently guided by a theoretical framework that can act as a supporting tool for the decision maker. For example, this may be applied to the case of vulnerability disclosure to estimate the effect in terms of the effective increase in risk of attacks that follows the disclosure, extending previous work in this same direction by Mitra and Ransbotham (2015). Our findings would indicate that there is a limited risk in additional disclosures for the same software version.

Finally, a more precise and data-grounded understanding of the attacker poses a strategic advantage for the defender. Extending Dey et al. (2015), parameters on attacker decisions and timings can be considered when choosing or evaluating a patching policy. Our findings contribute in this direction by defining a production model of attacks at scale that can be integrated in current evaluations; importantly, our empirical validation can serve as input for the parametrization of these models, guiding simulation scenarios and risk profile estimations to make better defensive decisions. For example, following Kannan et al. (2016), when an exogenous strategy is ‘imposed’ on an attacker (e.g. as defined by the release of malware in the underground markets), a non-uniform patch releases may be preferable from a vendor’s perspective. Our results may contribute in the definition of efficient policies whereby a desirable level of security is achieved (e.g. as defined by a social planner), while only marginally decreasing the profits for the vendor, by identifying which vulnerabilities can be ‘safely’ left unpatched on certain systems (e.g. with pirated software (Kannan et al., 2016; Terrence August, 2008)). Our results also suggest that the time dimension may be relevant to evaluate from a policy perspective, for example by asynchronously releasing patches to users. Similarly, by diversifying software (Chen et al., 2011; Homescu et al., 2013) the defender can effectively decrease the number of systems the attacker can compromise with one exploit, effectively making the existence conditions for Eq. (8) harder to satisfy than by means of a patch release strategy employed by vendors. Following this, software vendors may randomize the distribution of vulnerability patches to their users, to minimize the attacker’s chances of matching their exploits with the vulnerabilities actually present on the system. A random distribution of patches would simply decrease the fraction of attackable systems regardless of the attacker’s choice in which vulnerability to exploit. Moreover, diversifying defenses may be in fact less onerous than re-compiling code bases (when possible) or maintaining extremely diverse operational environments.

References

- Allodi, L. (2015). The heavy tails of vulnerability exploitation. In *Proceedings of the 2015 Engineering Secure Software and Systems Conference (ESSoS’15)*.
- Allodi, L., M. Corradin, and F. Massacci (2015). Then and now: On the maturity of the cybercrime markets the lesson that black-hat marketeers learned. *IEEE Transactions on Emerging Topics in Computing* 4(1), 35–46.
- Allodi, L., V. Kotov, and F. Massacci (2013). Malwarelab: Experimentation with cybercrime attack tools. In *Proceedings of the 2013 6th Workshop on Cybersecurity Security and Test*.
- Allodi, L. and F. Massacci (2014, August). Comparing vulnerability severity and exploits using case-control studies. *ACM Transaction on Information and System Security (TISSEC)* 17(1).
- Allodi, L. and F. Massacci (2017). Security events and vulnerability data for cybersecurity risk estimation. *Risk Analysis* 37(8), 1606–1627.

¹¹The NTIA forum can be found at <https://www.ntia.doc.gov/other-publication/2016/multistakeholder-process-cybersecurity-vulnerabilities>, last accessed January 15, 2017.

- Anderson, R. (2008). Information security economics - and beyond. In *Proceedings of the 9th International Conference on Deontic Logic in Computer Science, DEON '08*, pp. 49–49.
- Anderson, R. and T. Moore (2006). The economics of information security. *Science* 314.
- Arora, A., R. Krishnan, A. Nandkumar, R. Telang, and Y. Yang (2004). Impact of vulnerability disclosure and patch availability-an empirical analysis. In *Proceedings of the 3rd Workshop on Economics and Information Security*.
- Arora, A., R. Telang, and H. Xu (2008). Optimal policy for software vulnerability disclosure. *Management Science* 54(4), 642–656.
- Asghari, H., M. Van Eeten, A. Arnbak, and N. Van Eijk (2013). Security economics in the https value chain. In *Twelfth Workshop on the Economics of Information Security (WEIS 2013), Washington, DC*.
- August, T., M. F. Niculescu, and H. Shin (2014). Cloud implications on software network structure and security risks. *Information Systems Research* 25(3), 489–510.
- August, T. and T. I. Tunca (2006). Network software security and user incentives. *Management Science* 52(11), 1703–1720.
- August, T. and T. I. Tunca (2008). Let the pirates patch? an economic analysis of software security patch restrictions. *Information Systems Research* 19(1), 48–70.
- August, T. and T. I. Tunca (2011). Who should be responsible for software security? a comparative analysis of liability policies in network environments. *Management Science* 57(5), 934–959.
- Baltazar, J. (2011). More traffic, more money: Koobface draws more blood. Technical report, TrendLabs.
- Beattie, S., S. Arnold, C. Cowan, P. Wagle, C. Wright, and A. Shostack (2002). Timing the application of security patches for optimal uptime. In *LISA, Volume 2*, pp. 233–242.
- Bilge, L. and T. Dumitras (2012). Before we knew it: an empirical study of zero-day attacks in the real world. In *Proceedings of the 19th ACM Conference on Computer and Communications Security (CCS'12)*, pp. 833–844. ACM.
- Bilge, L., Y. Han, and M. Dell’Amico (2017). Riskteller: Predicting the risk of cyber incidents. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1299–1311. ACM.
- Birge, J. R. (2010). The persistence and effectiveness of large-scale mathematical programming strategies: projection, outer linearization, and inner linearization. In *A Long View of Research and Practice in Operations Research and Management Science*, pp. 23–33. Springer.
- Birge, J. R. and F. Louveaux (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Bozorgi, M., L. K. Saul, S. Savage, and G. M. Voelker (2010, July). Beyond heuristics: Learning to classify vulnerabilities and predict exploits. In *Proceedings of the 16th ACM International Conference on Knowledge Discovery and Data Mining*.
- Brand, M., C. Valli, and A. Woodward (2010). Malware forensics: Discovery of the intent of deception. *The Journal of Digital Forensics, Security and Law: JDFSL* 5(4), 31.
- Carlini, N. and D. Wagner (2014). Rop is still dangerous: Breaking modern defenses. In *23rd USENIX Security Symposium (USENIX Security 14)*, pp. 385–399.
- Cavusoglu, H., H. Cavusoglu, and J. Zhang (2006). Economics of security patch management. In *WEIS*.
- Chen, P.-y., G. Kataria, and R. Krishnan (2011). Correlated failures, diversification, and information security risk management. *MIS Quarterly* 35(2), 397–422.
- Christey, S. and B. Martin (2013, July). Buying into the bias: why vulnerability statistics suck. <https://www.blackhat.com/us-13/archives.html\#Martin>.
- CVSS-SIG (2015). Common vulnerability scoring system v3.0: Specification document. Technical report. First.org.
- DeGroot, M. H. (2005). *Optimal statistical decisions*, Volume 82. John Wiley & Sons.
- Dey, D., A. Lahiri, and G. Zhang (2015). Optimal policies for security patch management. *INFORMS Journal on Computing* 27(3), 462–477.
- Dolev, D. and A. Yao (1983a). On the security of public key protocols. *IEEE Transactions on information theory* 29(2), 198–208.
- Dolev, D. and A. Yao (1983b, mar). On the security of public key protocols. *IEEE Transactions on*

- Information Theory* 29(2), 198 – 208.
- Dumitras, T. and D. Shou (2011). Toward a standard benchmark for computer security research: The worldwide intelligence network environment (wine). In *Proceedings of the First Workshop on Building Analysis Datasets and Gathering Experience Returns for Security*, pp. 89–96. ACM.
- Erickson, J. (2008). *Hacking: the art of exploitation*. No Starch Press.
- Finifter, M., D. Akhawe, and D. Wagner (2013). An empirical study of vulnerability rewards programs. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, Washington, D.C., pp. 273–288. USENIX.
- Franklin, J., V. Paxson, A. Perrig, and S. Savage (2007). An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of the 14th ACM Conference on Computer and Communications Security (CCS'07)*, pp. 375–388.
- Frederick, S., G. Loewenstein, and T. O'donoghue (2002). Time discounting and time preference: A critical review. *Journal of economic literature* 40(2), 351–401.
- Frei, S., D. Schatzmann, B. Plattner, and B. Trammell (2010). Modeling the security ecosystem - the dynamics of (in)security. In T. Moore, D. Pym, and C. Ioannidis (Eds.), *Economics of Information Security and Privacy*, pp. 79–106. Springer US.
- Gal-Or, E. and A. Ghose (2005). The economic incentives for sharing security information. *Information Systems Research* 16(2), 186–208.
- Grier, C., L. Ballard, J. Caballero, N. Chachra, C. J. Dietrich, K. Levchenko, P. Mavrommatis, D. McCoy, A. Nappa, A. Pitsillidis, N. Provos, M. Z. Rafique, M. A. Rajab, C. Rossow, K. Thomas, V. Paxson, S. Savage, and G. M. Voelker (2012). Manufacturing compromise: the emergence of exploit-as-a-service. In *Proceedings of the 19th ACM Conference on Computer and Communications Security (CCS'12)*, pp. 821–832. ACM.
- Herley, C. (2013). When does targeting make sense for an attacker? *IEEE Security and Privacy* 11(2), 89–92.
- Herley, C. and D. Florencio (2010). Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. *Economics of Information Security and Privacy*.
- Homescu, A., S. Neisius, P. Larsen, S. Brunthaler, and M. Franz (2013). Profile-guided automated software diversity. In *2013 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pp. 1–11. IEEE.
- Houmb, S. H., V. N. Franqueira, and E. A. Engum (2010). Quantifying security risk level from cvss estimates of frequency and impact. *Journal of Systems and Software* 83(9), 1622–1634.
- Howard, M., J. Pincus, and J. Wing (2005). Measuring relative attack surfaces. *Computer Security in the 21st Century*, 109–137.
- Hutchings, A. and R. Clayton (2016). Exploring the provision of online booter services. *Deviant Behavior* 37(10), 1163–1178.
- Jonsson, E. and T. Olovsson (1997). A quantitative model of the security intrusion process based on attacker behavior. *IEEE Transactions on Software Engineering* 23(4), 235–245.
- Kanich, C., C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage (2008). Spamalytics: an empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS'08)*, CCS '08, pp. 3–14. ACM.
- Kannan, K., M. S. Rahman, and M. Tawarmalani (2016). Economic and policy implications of restricted patch distribution. *Management Science* 62(11), 3161–3182.
- Karthik Kannan, R. T. (2005). Market for software vulnerabilities? think again. *51*(5), 726–740.
- Kotov, V. and F. Massacci (2013). Anatomy of exploit kits. In J. Jürjens, B. Livshits, and R. Scandariato (Eds.), *Engineering Secure Software and Systems: 5th International Symposium, ESSoS 2013*, pp. 181–196. Springer Berlin Heidelberg.
- Laffont, J.-J. and D. Martimort (2009). *The theory of incentives: the principal-agent model*. Princeton University Press.
- Laszka, A., B. Johnson, and J. Grossklags (2013). Mitigation of targeted and non-targeted covert attacks as a timing game. In *International Conference on Decision and Game Theory for Security*, pp. 175–191. Springer.

- Lee, C. H., X. Geng, and S. Raghunathan (2016). Mandatory standards and organizational information security. *Information Systems Research* 27(1), 70–86.
- Liu, P., W. Zang, and M. Yu (2005, February). Incentive-based modeling and inference of attacker intent, objectives, and strategies. *ACM Transactions on Information and System Security* 8(1), 78–118.
- Manadhata, P. K. and J. M. Wing (2011). An attack surface metric. *IEEE Transactions on Software Engineering* 37, 371–386.
- Manshaei, M. H., Q. Zhu, T. Alpcan, T. Başçar, and J.-P. Hubaux (2013, July). Game theory meets network security and privacy. *ACM Comput. Surv.* 45(3), 25:1–25:39.
- Mell, P., K. Scarfone, and S. Romanosky (2007). A complete guide to the common vulnerability scoring system version 2.0. Technical report, FIRST, Available at <http://www.first.org/cvss>.
- Mellado, D., E. Fernández-Medina, and M. Piattini (2010). A comparison of software design security metrics. In *Proceedings of the Fourth European Conference on Software Architecture: Companion Volume, ECSA '10*, New York, NY, USA, pp. 236–242. ACM.
- Miller, C. (2007). The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Proceedings of the 6th Workshop on Economics and Information Security*.
- Mitra, S. and S. Ransbotham (2015). Information disclosure and the diffusion of information security attacks. *Information Systems Research* 26(3), 565–584.
- Murdoch, S. J., S. Drimer, R. Anderson, and M. Bond (2010). Chip and pin is broken. In *2010 IEEE Symposium on Security and Privacy*, pp. 433–446. IEEE.
- Naaliel, M., D. Joao, and M. Henrique (2014). Security benchmarks for web serving systems. In *Proceedings of the 25th IEEE International Symposium on Software Reliability Engineering (ISSRE'14)*.
- Nappa, A., R. Johnson, L. Bilge, J. Caballero, and T. Dumitras (2015). The attack of the clones: a study of the impact of shared code on vulnerability patching. In *2015 IEEE Symposium on Security and Privacy*, pp. 692–708. IEEE.
- Nayak, K., D. Marino, P. Efstathopoulos, and T. Dumitras (2014). Some vulnerabilities are different than others. In *Proceedings of the 17th International Symposium on Research in Attacks, Intrusions and Defenses*, pp. 426–446. Springer.
- Nguyen, V. H., S. Dashevskyi, and F. Massacci (2015). An automatic method for assessing the versions affected by a vulnerability. *Empirical Software Engineering*, 1–30.
- Nikiforakis, N., F. Maggi, G. Stringhini, M. Z. Rafique, W. Joosen, C. Kruegel, F. Piessens, G. Vigna, and S. Zanero (2014). Stranger danger: exploring the ecosystem of ad-based url shortening services. In *Proceedings of the 23rd international conference on World wide web*, pp. 51–62. ACM.
- NIST (2015). National institute of standards and technology, national vulnerability database. <http://nvd.nist.gov>.
- Okhravi, H. and D. Nicol (2008). Evaluation of patch management strategies. *International Journal of Computational Intelligence: Theory and Practice* 3(2), 109–117.
- Ooi, K. W., S. H. Kim, Q.-H. Wang, and K.-L. Hui (2012). Do hackers seek variety? an empirical analysis of website defacements. In *Proceedings of the 33rd International Conference on Information Systems (ICIS 2012)*. AIS.
- PCI-DSS (2010). PCI. https://www.pcisecuritystandards.org/documents/pci_dss_v2.pdf.
- Provos, N., P. Mavrommatis, M. A. Rajab, and F. Monroe (2008). All your iframes point to us. In *Proceedings of the 17th USENIX Security Symposium*, pp. 1–15.
- Quinn, S. D., K. A. Scarfone, M. Barrett, and C. S. Johnson (2010). Sp 800-117. guide to adopting and using the security content automation protocol (scap) version 1.0. Technical report, NIST.
- Ransbotham, S. and S. Mitra (2009). Choice and chance: A conceptual model of paths to information security compromise. *Information Systems Research* 20(1), 121–139.
- Ransbotham, S., S. Mitra, and J. Ramsey (2012). Are markets for vulnerabilities effective? *MIS Quarterly* 36(1), 43.
- Rao, J. M. and D. H. Reiley (2012, April). The economics of spam. *Journal of Economic Perspectives* 26(3), 87–110.
- Rincón-Zapatero, J. P. and C. Rodríguez-Palmero (2003). Existence and uniqueness of solutions to the bellman equation in the unbounded case. *Econometrica* 71(5), 1519–1555.

- Schneier, B. (2008). Inside the twisted mind of the security professional. *Wired Magazine*.
- Schwartz, E. J., T. Avgerinos, and D. Brumley (2011). Q: Exploit hardening made easy. In *USENIX Security Symposium*.
- Serra, E., S. Jajodia, A. Pugliese, A. Rullo, and V. Subrahmanian (2015a). Pareto-optimal adversarial defense of enterprise systems. *ACM Transactions on Information and System Security (TISSEC)* 17(3), 11.
- Serra, E., S. Jajodia, A. Pugliese, A. Rullo, and V. S. Subrahmanian (2015b, March). Pareto-optimal adversarial defense of enterprise systems. *ACM Trans. Inf. Syst. Secur.* 17(3), 11:1–11:39.
- Sheyner, O., J. Haines, S. Jha, R. Lippmann, and J. M. Wing (2002). Automated generation and analysis of attack graphs. *Proceedings of the 29th IEEE Symposium on Security and Privacy* 0, 273.
- Smeets, M. (2018). A matter of time: On the transitory nature of cyberweapons. *Journal of Strategic Studies* 41(1-2), 6–32.
- Stock, B., S. Lekies, and M. Johns (2013). 25 million flows later-large-scale detection of dom-based xss. *20th CCS. ACM*.
- Stokey, N. L. (2008). *The Economics of Inaction: Stochastic Control models with fixed costs*. Princeton University Press.
- Stone-Gross, B., M. Cova, L. Cavallaro, B. Gilbert, M. Szydowski, R. Kemmerer, C. Kruegel, and G. Vigna (2009). Your botnet is my botnet: analysis of a botnet takeover. In *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS'09)*. ACM.
- Telang, R. and S. Wattal (2007). An empirical analysis of the impact of software vulnerability announcements on firm stock price. *Software Engineering, IEEE Transactions on* 33(8), 544–557.
- Terrence August, T. I. T. (2008). Let the pirates patch? an economic analysis of software security patch restrictions. *Information Systems Research* 19(1), 48–70.
- van Dijk, M., A. Juels, A. Oprea, and R. L. Rivest (2013, Oct). Flipit: The game of “stealthy takeover”. *Journal of Cryptology* 26(4), 655–713.
- Van Eeten, M. and J. Bauer (2008). Economics of malware: Security decisions, incentives and externalities. Technical report, OECD.
- Van Eeten, M., J. Bauer, H. Asghari, S. Tabatabaie, and D. Rand (2010). The role of internet service providers in botnet mitigation: An empirical analysis based on spam data. Technical report, OECD STI Working Paper.
- Verizon (2011). 2011 data breach investigation report. Technical report, Verizon.
- Wang, J. A., H. Wang, M. Guo, and M. Xia (2009). Security metrics for software systems. In *Proceedings of the 47th Annual Southeast Regional Conference, ACM-SE 47*, New York, NY, USA, pp. 47:1–47:6. ACM.
- Wang, L., T. Islam, T. Long, A. Singhal, and S. Jajodia (2008). An attack graph-based probabilistic security metric. In *Proceedings of the 22nd IFIP WG 11.3 Working Conference on Data and Applications Security*, Volume 5094 of *Lecture Notes in Computer Science*, pp. 283–296. Springer Berlin / Heidelberg.
- Wash, R. (2010). Folk models of home computer security. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*.
- Yeo, M. L., E. Rolland, J. R. Ulmer, and R. A. Patterson (2014). Risk mitigation decisions for IT security. *ACM Transactions on Management Information Systems (TMIS)* 5(1), 5.
- Zhao, M., A. Laszka, and J. Grossklags (2017). Devising effective policies for bug-bounty platforms and security vulnerability discovery. *Journal of Information Policy* 7, 372–418.
- Zhuge, J., T. Holz, C. Song, J. Guo, X. Han, and W. Zou (2009). *Studying Malicious Websites and the Underground Economy on the Chinese Web*, pp. 225–244. Boston, MA: Springer US.

7 Appendix: Proofs of Propositions

7.1 Proof of Proposition 1

The objective of the Proposition is to demonstrate the solution condition for the optimal set of action times $\{T_1^*, \dots, T_n^*\}$, which is given by the recursive derivative where $T_0 = 0$, $T_{n+1} = \infty$, $\delta, \lambda > 0$.

$$\frac{\partial \Pi(T_i, T_{i-1})}{\partial T_i} e^{-\delta T_{i-1}} - \delta(\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i} + \frac{\partial \Pi(T_{i+1}, T_i)}{\partial T_i} e^{-\delta T_i} = 0 \quad (16)$$

Proof. Proof of Proposition 1. We solve Eq. (3) in expectations and replace Eq. (4) and $c_i(t) \approx 0$ in Eq. (3) with the abbreviations $\theta_{-1} \equiv 0$, $\theta_0 \equiv \theta_V$, and $\theta_i \equiv \theta_{V \cup \{v_1 \dots v_i\}}$. Hence, we obtain

$$\{T_1^*, \dots, T_n^*\} \approx \arg \max_{\{T_1, \dots, T_n\}} \sum_{i=0}^n -C_i e^{-\delta T_i} + \int_{T_i}^{T_{i+1}} rN \left(\theta_{i-1} e^{-\lambda t} + (\theta_i - \theta_{i-1}) e^{-\lambda(t-T_i)} \right) e^{-\delta t} dt \quad (17)$$

We can now solve the integral by replacing $t \rightarrow z + T_i$

$$\begin{aligned} & \int_0^{T_{i+1}-T_i} rN \left(\theta_{i-1} e^{-\lambda(z+T_i)} + (\theta_i - \theta_{i-1}) e^{-\lambda(z)} \right) e^{-\delta(z+T_i)} dz = \dots \\ & = rN e^{-\delta T_i} (\theta_i + \theta_{i-1} e^{-\lambda T_i} - \theta_{i-1}) \int_0^{T_{i+1}-T_i} e^{-(\delta+\lambda)z} dz \\ & = \frac{rN}{\delta+\lambda} e^{-\delta T_i} (\theta_i - \theta_{i-1} + \theta_{i-1} e^{-\lambda T_i}) \left(1 - e^{-(\delta+\lambda)(T_{i+1}-T_i)} \right) \end{aligned} \quad (18)$$

Hence we finally obtain the following result which can be rewritten as (5)

$$\{T_1^*, \dots, T_n^*\} = \arg \max_{\{T_1, \dots, T_n\}} \sum_{i=0}^n e^{-\delta T_i} \left(-C_i + \frac{rN}{\lambda+\delta} (\theta_i - \theta_{i-1} + \theta_{i-1} e^{-\lambda T_i}) \left(1 - e^{-(\lambda+\delta)(T_{i+1}-T_i)} \right) \right) \quad (19)$$

To identify the optimal T_i we take the usual first order condition and obtain for $i = 1 \dots n$

$$\begin{aligned} \frac{\partial \Pi}{\partial T_i} &= \frac{\partial}{\partial T_i} (\dots + (\Pi(T_{i-1+1}, T_{i-1}) - C_{i-1}) e^{-\delta T_{i-1}} + (\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i} \\ &\quad + (\Pi(T_{i+1+1}, T_{i+1}) - C_{i+1}) e^{-\delta T_{i+1}} + \dots) \\ &= \dots + \frac{\partial}{\partial T_i} ((\Pi(T_i, T_{i-1}) - C_{i-1}) e^{-\delta T_{i-1}} + (\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i}) + \dots \\ &= \frac{\partial \Pi(T_i, T_{i-1})}{\partial T_i} e^{-\delta T_{i-1}} - \delta(\Pi(T_{i+1}, T_i) - C_i) e^{-\delta T_i} + \frac{\partial \Pi(T_{i+1}, T_i)}{\partial T_i} e^{-\delta T_i} \end{aligned} \quad (20)$$

□

7.2 Proof of Corollary 1

Corollary 1 outlines the solution space when the attacker presumes the residual income after T_1^* is fixed in time t expectations, hence the attacker is myopic to action time T_1^* .

Proof. Proof of Corollary 1 For $n = 1$ Eq. (7) can be simplified by substituting $T_0 = 0$ and $T_{n+1} = \infty$ and $T_n = T$.

$$\frac{\partial \Pi(T, 0)}{\partial T} - \delta(\Pi(\infty, T) - C_V |V) e^{-\delta T} + \frac{\partial \Pi(\infty, T)}{\partial T} e^{-\delta T} = 0 \quad (21)$$

To determine the three components of the equation above we now decompose each of the individual terms of Eq. (6) as follows:

$$\Pi(T, 0) = \frac{rN}{\lambda+\delta} \theta_V \left(1 - e^{-(\lambda+\delta)T} \right) \quad \text{and} \quad \Pi(\infty, T) = \frac{rN}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V + \theta_V e^{-\lambda T}) \quad (22)$$

we can now derive the partial derivatives

$$\frac{\partial \Pi(T,0)}{\partial T} = rN\theta_V e^{-(\lambda+\delta)T} \quad \text{and} \quad \frac{\partial \Pi(\infty,T)}{\partial T} = -rN \frac{\lambda}{\lambda+\delta} \theta_V e^{-\lambda T} \quad (23)$$

We then replace the corresponding value in the Eq. (21) above:

$$\frac{\partial \Pi}{\partial T} = rN\theta_V e^{-(\lambda+\delta)T} - \delta \left(\frac{rN}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V + \theta_V e^{-\lambda T}) - C(v|V) \right) e^{-\delta T} - rN \frac{\lambda}{\lambda+\delta} \theta_V e^{-\lambda T} e^{-\delta T} \quad (24)$$

$$= \frac{rN}{\lambda+\delta} e^{-\delta T} \left((\lambda + \delta) \theta_V e^{-\lambda T} - \delta \left((\theta_{V \cup \{v\}} - \theta_V + \theta_V e^{-\lambda T}) - (\lambda + \delta) \frac{C(v|V)}{rN} \right) - \lambda \theta_V e^{-\lambda T} \right) \quad (25)$$

which finally yields

$$\partial \Pi / \partial T = rN e^{-\delta T} \left(\frac{C(v|V)}{rN} - \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) \right) \quad (26)$$

Now observe that if $\frac{C(v|V)}{rN} > \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V)$ then the derivative is positive decreasing and it is convenient to postpone the update which is eventually reached for $T^* \rightarrow \infty$. This is particularly true when $\theta_{V \cup \{v\}} - \theta_V = 0$ and namely there is no change in the number of infected systems by adding one more vulnerability. If $\frac{C(v|V)}{rN} \leq \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V)$ the derivative is negative so any update would decrease the marginal return. Only if $\frac{C(v|V)}{rN} = \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V)$ then the derivative is identically zero and the attacker is indifferent to the time of deployment. \square

7.3 Proof of Proposition 2

The final theoretical result provides an explicit solution to the myopic version of the dynamic programming problem and proceeds as follows:

Proof. We impose the stopping condition to the f.o. derivative of the profit of the attacker in Eq. (26)

$$\partial \Pi / \partial T = rN e^{-\delta T} \left(\frac{C(v|V)}{rN} - \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) \right) = \frac{r(0, N_V(0), V)}{N_V(0)} = r \quad (27)$$

$$N \left(\frac{C(v|V)}{rN} - \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) \right) = e^{\delta T} \quad (28)$$

$$T_r = \frac{1}{\delta} \log \left(\frac{C(v|V)}{r} - \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) N \right) \quad (29)$$

As the exploit weaponization has to happen for $T_r \geq 0$ we must have $\frac{C(v|V)}{r} - \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) N \geq 1$ and therefore $C(v|V) \geq r + \frac{\delta}{\lambda+\delta} (\theta_{V \cup \{v\}} - \theta_V) rN$. \square

7.4 Proof of Claim 1

The transformation of the model prediction to the number of attacks against $\theta_{V \cup \{v\}} N$ systems received \mathcal{T} days after receiving an attack against a different vulnerability is as follows.

Proof. Setting the number of attacks on the first vulnerability at time $t - \mathcal{T}$ as $\mathcal{U}(\theta_v, t - \mathcal{T}) = N\theta_v e^{-\lambda(t-\mathcal{T})}$ and the attacks received on the second vulnerability at time t as $\mathcal{U}(\theta_{V \cup \{v\}}, t) = N\theta_{V \cup \{v\}} e^{-\lambda(t-T^*)}$, we obtain that the expected attacks received \mathcal{T} days after the first attack are as follows:

$$\begin{aligned} \mathcal{U}(\theta_{V \cup \{v\}}, t, \mathcal{T}) &= \int_{\max(\mathcal{T}, T^*)}^{+\infty} \min \left(N\theta_v e^{-\lambda(t-\mathcal{T})}, N\theta_{V \cup \{v\}} e^{-\lambda(t-T^*)} \right) dt \\ &= \min(N\theta_v e^{\lambda \mathcal{T}}, N\theta_{V \cup \{v\}} e^{\lambda T^*}) \int_{\max(\mathcal{T}, T^*)}^{+\infty} e^{-\lambda t} dt \\ &= \frac{1}{\lambda} \min(N\theta_v e^{\lambda \mathcal{T}}, N\theta_{V \cup \{v\}} e^{\lambda T^*}) e^{-\lambda(\max(\mathcal{T}, T^*))} \end{aligned} \quad (30)$$

$$\log \mathcal{U}(\theta_{V \cup \{v\}}, t, \mathcal{T}) = \log \frac{1}{\lambda} + \min(\log N\theta_v + \lambda \mathcal{T}, \log N\theta_{V \cup \{v\}} + \lambda T^*) - \lambda(\max(\mathcal{T}, T^*)) \quad (31)$$

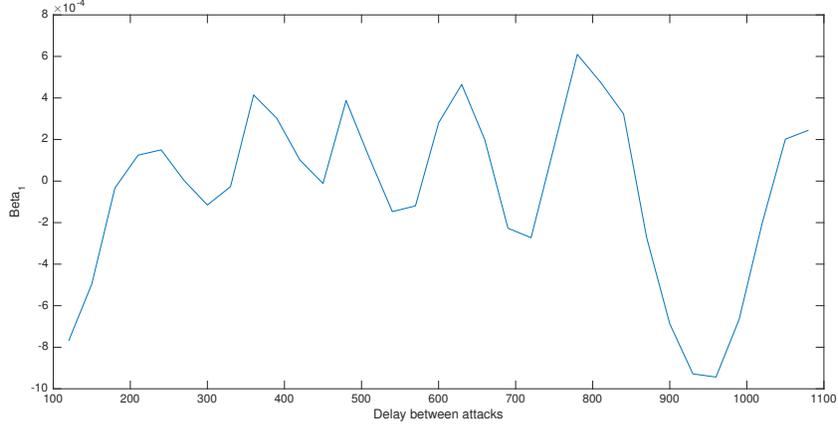


Figure 8: β_1 estimations for increasing values of \mathcal{T} .

Solve for the case $T^* > \mathcal{T}$. As $N\theta_V \leq N\theta_{V \cup \{v\}}$, we have that $\min(\log N\theta_V + \lambda\mathcal{T}, \log N\theta_{V \cup \{v\}} + \lambda T^*) = \log N\theta_V + \lambda\mathcal{T}$, and we obtain:

$$T^* > \mathcal{T}, \quad \log \mathcal{U} = \log \frac{N}{\lambda} + \log \theta_V + \lambda\mathcal{T} - \lambda T^* \quad (32)$$

Solve for the case $T^* \leq \mathcal{T}$. Sub case 1. For $\log N\theta_V + \lambda\mathcal{T} \leq \log N\theta_{V \cup \{v\}} + \lambda T^*$ we obtain:

$$\{T^* \leq \mathcal{T} \leq \frac{1}{\lambda} \log \frac{\theta_{V \cup \{v\}}}{\theta_v} + T^*, \quad \log \mathcal{U} = \log \frac{N\theta_V}{\lambda}\} \quad (33)$$

which indicates that, within a small timeframe after the introduction of the exploit at time T^* , the number of received attacks only depends on the number of vulnerable systems in the wild. This result appears to explain the observation noted in Mitra and Ransbotham (2015).

Sub case 2. For $\log N\theta_V + \lambda\mathcal{T} > \log N\theta_{V \cup \{v\}} + \lambda T^*$ we obtain:

$$\{\mathcal{T} > \frac{1}{\lambda} \log \frac{\theta_{V \cup \{v\}}}{\theta_V} + T^*, \quad \log \mathcal{U} = \log \frac{N}{\lambda} + \log \theta_V + \log \theta_{V \cup \{v\}} - \log \theta_V + \lambda T^* - \lambda\mathcal{T}\} \quad (34)$$

By comparing (32) with (34) it is apparent that the impact of \mathcal{T} on $\log \mathcal{U}$ should change in sign wrt \mathcal{T} . \square

7.5 Sensitivity Analysis

Figure 8 plots the estimated coefficient β_1 for the empirical variable \mathcal{T} by constraining intervals of 120 days (four months) for increasing values of \mathcal{T} (e.g. $0 < \mathcal{T} \leq 120$, $30 < \mathcal{T} \leq 150$, $60 < \mathcal{T} \leq 180$, ...). It can be observed that the sign of β_1 oscillates above and below zero with increasing amplitude as \mathcal{T} increases. This effect is present regardless of the size of the interval and the relative increment imposed on \mathcal{T} .

A On-line Appendix: Internet Explorer Exploit Code

Here we report the exploit code for $CVE_1 = CVE-2010-0806$ and $CVE_2 = CVE-2009-3672$ both affecting Internet Explorer version 7. On Exploit-DB¹² CVE_1 corresponds to exploit 16547 and CVE_2 corresponds to exploit 11683.

The exploit code fragments (10+ lines of code out of 260+ for CVE_1 and 130+ for CVE_2) below illustrates the difference between the two exploits as scripted for the Metasploit engine. The exploit code for CVE_2 is effectively a rearrangement of the exploit code for CVE_1 , with different variable names (e.g. by replacing `j_memory` with `var_memory`, `j_shellcode` with `var_shellcode`, etc.) and repositioned at the appropriate memory addresses (0x...).

```
1 function j_function1 () {
2   ...
3   j_memory = new Array ();
4   var j_shellcode = unescape (...);
5   var j_slackspace = 0x86000 -(j_shellcode.length*2);
6   while(j_nops.length < j_slackspace/2)
7     j_nops+=j_nops;
8   for(j_counter=0; j_counter < 270; j_counter++) {
9     j_memory[j_counter] = j_fillblock + j_fillblock + j_shellcode; }}
```

```
1 function var_body(){
2   ...
3   var var_memory = new Array ();
4   var var_shellcode = var_unescape (...);
5   var var_ss = 20 + var_shellcode.length;
6   while (var_spray.length < var_ss)
7     var_spray+=var_spray;
8   var_bk = var_spray.substring (...);
9   while(var_bk.length+var_ss < 0x100000)
10    var_bk = var_bk + var_bk + var_fb;
11   for (var_i=0; var_i < 1285; var_i++) {
12    var_memory[var_i]= var_bk + var_shellcode; }}
```

B On-line Appendix: Replication Guide

Here we give the replication guidelines for our study from the tables from WINE. Our data can be reproduced by utilizing the reference data set *WINE-2012-008*, archived in the WINE infrastructure.

Basic Tables from WINE

The first table we construct from WINE, LIFE, reports the full history of attacks against distinct users in WINE. It has the following structure:

	UserID	AttackSignature	Date	SystemID	IP_HASH	VolumeAttacks
LIFE =	:	:	:	:	:	:

¹²See Exploit-DB (<http://www.exploit-db.com>, last accessed April 12, 2018.)

Table 10: Number of WINE users in the **Hst**, **Frq**, and **Pk** control groups.

Hst	#Users	Frq	#Users	Pk	#Users
STABLE	1,446,020	LOW	3,210,465	LOW	2,559,819
ROAM	96589	MEDIUM	311,929	MEDIUM	983,221
UPGRADE	306,856	HIGH	19,683	HIGH	3,783
EVOLVE	1,697,570	VERYHIGH	3,919	VERYHIGH	170
		EXTREME	1,039	EXTREME	42

where $SID_1=SignatureID$ and $SID_2=SignatureID_later$ identify respectively the attack signature triggered in the first and the second attack received by the user. \mathcal{T} reports the days passed in between the two attacks. $\mathcal{U}=\text{count}(\text{machineID})$ reports the number of machines affected by $SignatureID_later$ \mathcal{T} days after receiving an attack of type $SignatureID$. Similarly, $\mathcal{N}=\text{sum}(\text{volumes of } SignatureID_later)$ counts how many times these two attacks triggered an alarm \mathcal{T} days apart. The difference with the previous field is that single machines may receive more than one attack of this type, thus we have $\text{sum}(\text{volumes of } SignatureID_later) \geq \text{count}(\text{machineID})$. The remaining fields are obtained with a join with the tables `PLATFORM`, `USERS`, `TARGET_PROFILE`.

Merging WINE with CVEs

`SignatureID` and `SignatureID_later` are internal codes that identify which attack signature deployed by Symantec’s product the attack triggered. To identify which vulnerability (if any) the attack attempted to exploit we map WINE’s `SignatureID` with the threat description publicly available at Symantec’s Security Response dataset.¹³ In the attack description it is provided, when relevant, the vulnerability that the attack exploits, as referenced by the unique CVE vulnerability identifier.¹⁴ The CVE-ID is a standard reference identifier for software vulnerabilities introduced by the MITRE organization and used by all major vulnerability databases such as the National Vulnerability Database, NVD.¹⁵

To characterize each vulnerability, we match the CVE-ID reported in the `SignatureID` with the vulnerability summary reported in the NVD. The information on NVD comprises the name of the affected software `Sw` (e.g. Flash in the example above), the latest vulnerable version `Ver` of the software (in our example Flash 9.0.115 and 9.0.45), and the disclosure date `Day`.

Further, additional information describing the technical aspects of the vulnerability are also provided. This information can be extracted from the Common Vulnerability Scoring System (CVSS) assessment of the vulnerability. CVSS measures several technical dimensions of a vulnerability to obtain a standardized assessment that can be used to meaningfully compare software vulnerabilities. However, previous studies showed that not all measures show, in practice, enough variability to characterize the vulnerabilities Allodi and Massacci (2014). Of the dimensions considered in CVSS, in this study we are specifically interested in the `Access Complexity` and `Imp` measures. The former gives an assessment on the ‘difficulty’ associated with engineering a reliable exploit for the vulnerability Mell et al. (2007). For example, a vulnerability that requires the attacker to win a race condition on the affected system in order to successfully exploit kit may be deemed as a High complexity vulnerability (because the attacker can not directly control the race condition, thus exploitation can be only stochastically successful). Similarly, `Imp` gives an assessment on the Confidentiality, Integrity and Availability losses that may follow the exploitation of the vulnerability.

Table 10 reports the count distribution of the number of WINE users for each of these factor’s levels. It is apparent that in the case of `Hst` users are uniformly distributed among the factor levels, with `ROAM` users being the least frequent category. Most of the mass of the `Frq` and `Pk` distributions is at the low end of the scale (i.e. most users receive few attacks per day both as an average and as a maximum). From Table 10 it

¹³The reference dataset is at <https://www.symantec.com/security-center/>, last accessed April 12, 2018.

¹⁴The classifiers are available at <http://cve.mitre.org>, last accessed April 12, 2018.

¹⁵Full database can be found at <http://nvd.nist.gov>, last accessed April 12, 2018.

appears that users characterized by EXTREME or VERYHIGH levels for Frq or Pk are outliers and may therefore need to be controlled for.